“Turnkey” Approach to Statewide Data Center

WA State Department of Information Services

July 2, 2010

Submitted by

Robin Currin
Regional Public Sector Manager
Ph: 360.584.4688
Fax: 360.252.7763
E-mail: RobinC@INXi.com

John Chesney
State of WA Account Manager
Ph: (206) 412-6504
E-mail: JohhC@INXi.com

INX, Inc.
500-108th Ave. NE,
Suite #1400
Bellevue, WA 98004

February 23, 2010

Rebekah O'Hara
Department of Information Services
Office of Legal Services
Department of Information Services
1110 Jefferson Street SE
PO Box 42445
Olympia, WA 98504

rebekaho@dis.wa.gov

RE: "Turnkey" Approach to State Wide Data Center

Dear Ms. O'Hara:

INX would like to thank the State of Washington Department of Information Services for the opportunity to bid on the Turnkey" Approach to State Wide Data Center**.**

The RFI response consists of the following materials, documents and enclosures:

- Cover Letter
- RFI Response
- Proposed Staff Vitae
- INX Virtual Hosting/Recovery Service Brochures

INX is a unique solution provider; we align ourselves with only the top manufacturing partners and commit to the technology by investing what it takes to achieve the highest level of certification possible with each of our vendor partners.

INX, Inc. (NASDAQ: INXI) is a professional consulting organization, purpose built to design and implement advanced technology solutions in Information Lifecycle with industry leaders in the enterprise Data Management and IP Communications industry. These enterprise leaders include: Cisco Systems, VMware, Microsoft, Citrix, EMC, NetApp, IBM, HP, Dell, Sun, Oracle, SAP and several 3rd party application providers. INX is a unique solution provider in that we align ourselves with only the top manufacturing partners and commit to the technology by investing in our resources (engineering, lab, test & development, etc.) to achieve the highest level of certification possible with each of our partners.

Sincerely,

Robin Currin

Robin Currin
Public Sector Manager
(360) 584-4688

Table of Contents

Briefly describe the essential characteristics of your approach to this "turnkey" strategy to the statewide data center. Highlight elements of the proposed model that would make it particularly attractive to the state and elements that will make it particularly difficult to implement.

Turnkey Approach Overview:

INX proposes the state of Washington adopt an approach that allows cost savings through near term agency virtualization efforts with the future end state in mind. As Washington state agencies virtualize incrementally, prior to SDC occupancy, the virtual machines and application data must remain on a migration path to a containerized shared services cloud. Keeping initial server consolidation on the roadmap to the State Data Center (SDC) cloud requires virtual storage replication strategies as a critical component of the interim design. Long term consolidation cost benefits will not be realized if the roadmap to the end state infrastructure is ignored. Our proposed end state cloud approach is based on the principles we used for building, scaling, managing, governing, and optimizing our own hosting business model within the SDC: Virtual Hosting Service (VHS).

INX's turnkey approach is a holistic approach that considers near term server consolidation and end state cloud services.

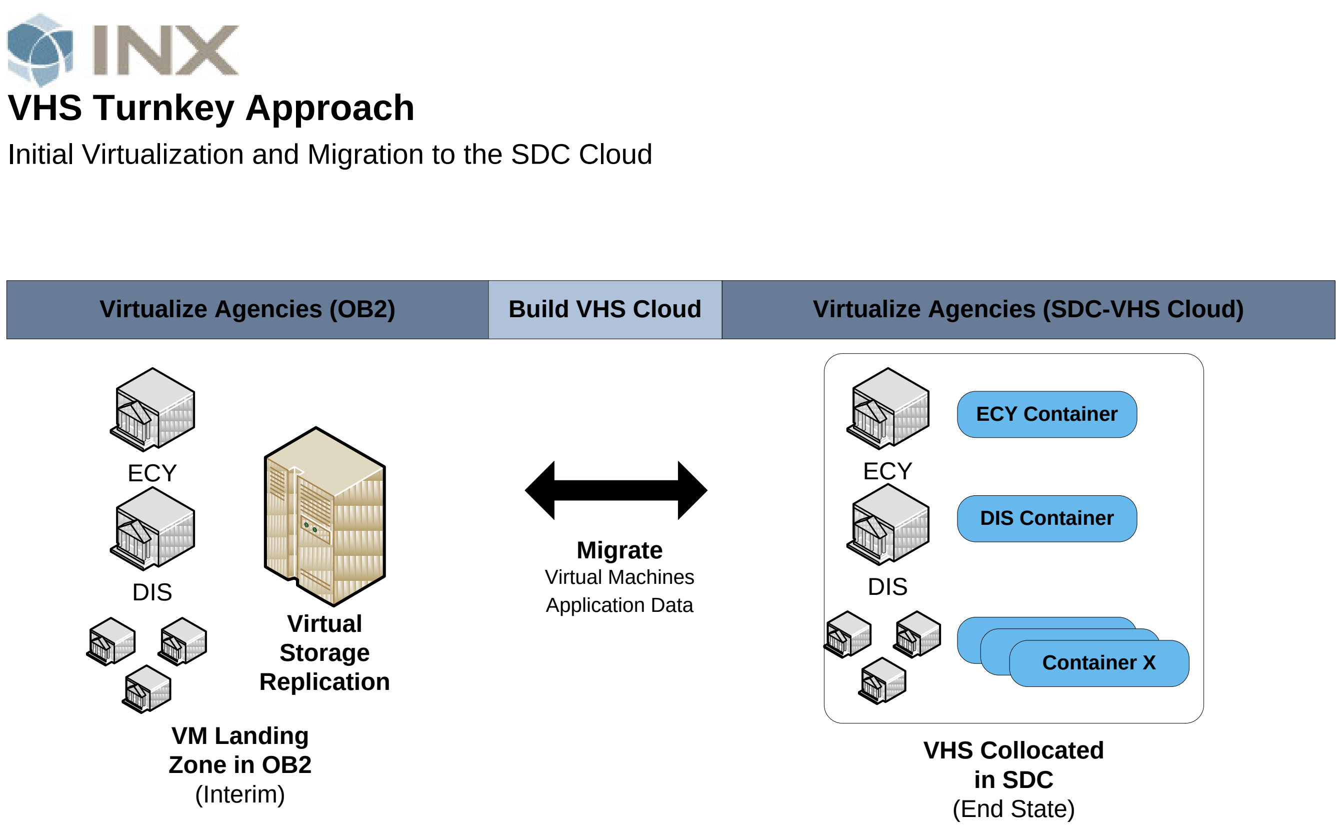

Once DIS gains occupancy of the SDC, INX suggests we build, manage, operate and transition a VHS cloud infrastructure collocated in the SDC. Once the first set of SDC-VHS containers are production ready, INX will guide DIS and their virtualized customers through a migration process to the new shared services cloud. At a high level, the migration process will move the virtual machines and application data over-the-wire to the SDC. Since OB2 will be the first datacenter to make the transition to the SDC, it is a logical conclusion to have OB2 be the target 'landing zone' for early centralized server consolidation. The transition from OB2 to SDC can be done with vMotion and virtual data replication technologies to assure minimal impact to state business.

VHS Overview

Building on INX's knowledge and expertise in virtualization, networking, storage and the data center, the Virtual Infrastructure Hosting Service ® (VHS) delivers unsurpassed highly-available capabilities utilizing VMware ESX and other VMware products, monitoring and management software, multiple, fiber channel SANs, a combination of the industry's newest 4-processor, quad core servers, Gig-E switches, firewall, load balancing and other components. It includes 24 X 7 monitoring and management along with BCP connections to multiple fiber carriers. The service is hosted at one of INX's multiple, top tier data centers, with minimum N+1 redundancy in all electrical, cooling and fiber network systems – data centers designed, built and operated for 99.999% availability.

INX's onsite engineers and technicians are fully trained and experienced in all aspects of the VHS hardware and software. Using VMware's Virtual Infrastructure as the foundation, VHS is architected, built and managed to achieve 99.95% system availability including both - planned and unplanned downtime. INX engineers work closely with clients to quickly handle problems that occur, i.e., such as restoring accidently deleted files, adding additional memory or storage capacity when needed.

Virtualized Disaster Recovery

Using INX's Virtual Infrastructure Hosting Service ®™ (VHS), allows for the simple implementation of INX's Virtual Infrastructure Recovery Service ®™ (VRS) which enables failover of all X86 servers to a second data center in the event of a disaster. A virtual replica of your production environment is up and running in less than 4 hours! Because we virtualize client access, we enable continued computer operations in the event of a prolonged pandemic disaster.

Benefits of INX Virtual Hosting Service (VHS) approach:

- Allows the state of WA to invest incrementally as the virtual infrastructure scales
- Manages risk through incremental investment (inflection point of 300 servers)
- Uses a cost effective monthly 'phone bill' chargeback mechanism for departmental tenants
- No upfront or costly capital expenses for servers, SAN storage and network infrastructure
- Minimizes costs of capital depreciating assets for the state of WA during the hosting period
- Reduced operating costs for systems administration and systems maintenance
- INX architects and engineer employ industry best practices to design and operate the VHS
- INX periodically upgrades the VHS – so you are operating your business on the newest and fastest computing infrastructure

How Does VHS Work?

Building on INX's knowledge and expertise in virtualization, networking, storage and the data center, the Virtual Infrastructure Hosting Service ® (VHS) delivers unsurpassed highly-available capabilities utilizing VMware ESX and other VMware products, monitoring and management software, multiple, fiber channel SANs, a combination of the industry's newest 4-processor, quad core servers, Gig-E switches, firewall, load balancing and other components. It includes 24 X 7 monitoring and management along with BCP connections to multiple fiber carriers. The service is hosted at one of INX's multiple, top tier data centers, with minimum N+1 redundancy in all electrical, cooling and fiber network systems – data centers designed, built and operated for 99.999% availability.

INX's onsite engineers and technicians are fully trained and experienced in all aspects of the VHS hardware and software. Using VMware's Virtual Infrastructure as the foundation, VHS is architected, built and managed to achieve 99.95% system availability including both - planned and unplanned downtime. INX engineers work closely with clients to quickly handle problems that occur, i.e., such as restoring accidently deleted files, adding additional memory or storage capacity when needed.

Virtualized Disaster Recovery

Using INX's Virtual Infrastructure Hosting Service ®™ (VHS), allows for the simple implementation of INX's Virtual Infrastructure Recovery Service ®™ (VRS) which enables failover of all X86 servers to a second data center in the event of a disaster. A virtual replica of your production environment is up and running in less than 4 hours! Because we virtualize client access, we enable continued computer operations in the event of a prolonged pandemic disaster.

Cost

Cost for the Virtual Infrastructure Hosting Service® depends upon the specific configuration and includes significant volume discounts. The price includes the cost savings normally associated with the server consolidation benefits of virtualization. The ROI for using VHS provides a compelling case for the use of this great service. Refer to Section 6 for more cost per virtual unit detail.

Self-provisioning portal

In large data center environments, and especially in those that are cloud-enabled, it is a well known fact, that a large amount of Virtual Machines in the infrastructure may heavily tax the operational readiness aspects of IT management.

The provisioning process to comply with the demand of VM requests becomes a tactical nightmare, and the overall efficiency of the data center begins to decline. The industry recognizes this syndrome with the term "Virtualization Sprawl".

Elements that might make this approach difficult to implement:

The three major operational-readiness-related challenges that come to light are:

1. Lack of change control process
    a. VMs being provisioned without proper approval
    b. Process for approving requests
    c. Audit trail on changes and by whom

2. Lack of automation and intelligence
    a. Manual steps…inconsistency or error
    b. VM placement with respect to resources
    c. How to translate business needs to technical
3. Lack of visibility
    a. VM ownership
    b. VM usage and need
    c. VM decommissioning

In order to avoid falling into these challenges, we envision the implementation of a Web Portal capable of providing a consistent automated process for managing the lifecycles of virtual machines.
This web interface will manage the interactions among everyone involved in the lifecycle of a virtual machine, from provisioning to decommissioning. During setup, IT administrators will create a catalog of virtual machine templates for the users to view and select from. The templates will help users determine the characteristics of available virtual machines (machine size, memory, storage, backup services, etc.). IT administrators will also define what types of approvals are required prior to virtual machine deployment.

After the initial setup is completed, users will then log into the portal, and submit requests for virtual machines. During the request process, users will specify the service template and quality of service they need. Based on these request criteria the portal will determine the specific resources that best support the request. For example, a request for a production virtual machine requiring high performance will be mapped to the highest performing server, network and storage resources available in the virtualized environment.

Once submitted, requests will be forwarded to an IT administrator for review and approval. If approved, automated virtual machine provisioning and placement will occur, based on previously defined lifecycle management policies. Users will then receive notification via email when their virtual machines are deployed and accessible.

When a virtual machine reaches the end of its lifecycle, the portal will decommission it. The decommission process, which will consist of an optional archiving step, and ultimately deleting the virtual machine, will provide better resource utilization, by ensuring resources come back into the resource pool for future use. If needed, users will be able to request a lifecycle extension that will then be reviewed and approved by an IT administrator, instead of letting the virtual machine be decommissioned.

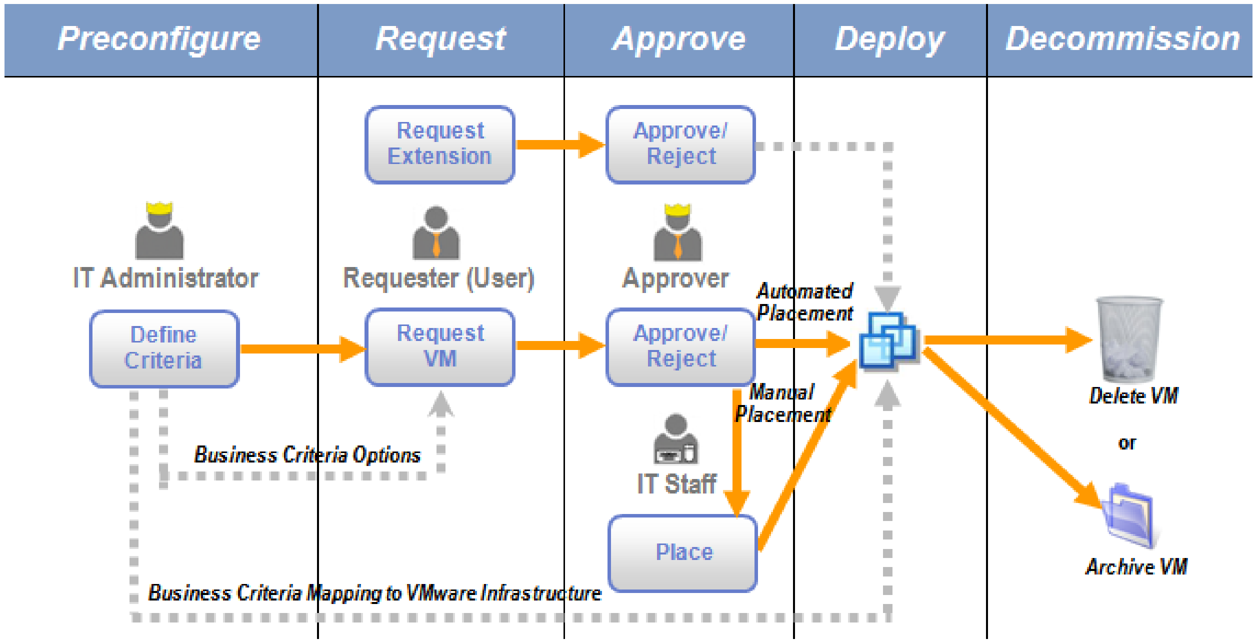

In a nut shell, we envision this portal the as a preventive and proactive measure to mitigate any virtualization sprawl that may occur, by incorporating the following three major operational-readiness-related functions:

1. Service catalog:
    - To allow IT staff to pre-define virtual machine configurations as a Web-based service to standardize virtual machine provisioning requests.
    - To standardize virtual machine configurations according to IT policies.
    - To create and enable service tiers that will define what quality of service levels get configured during virtual machine deployment (availability, security, monitoring).
2. Provisioning workflow automation:
    - To automate the routine provisioning management tasks, reclaim unused resources and enforce IT policies across virtual machine deployments.
    - To automate virtual machine guest customizations and system or network configuration tasks based on pre-defined policies.
    - To automate resource reclamation and automatically free up resources by decommissioning virtual machines that have reached the end of their lifecycle.
3. Administrative Control:
    - To provide IT staff with a Web-based portal that will centralize virtual machine requests and streamline the approval process. Through the portal, users will also be able to establish the service catalog, configure resource mappings and define placement policies.
    - To easily track the status of virtual machines through request, approval, provisioning and decommissioning.

Section 2 – Responses to Questions regarding the Overview

Briefly address each of the four (4) questions in the Overview section of the RFI description above. The questions represent threshold issues: is this approach worthy of further consideration? What issues might be relevant to deciding whether to put further effort into pursuing this direction?

Question 1: What are the financial, legal, and operational advantages and disadvantages of the model under investigation relative to a fully state run transition and operation? Are there unique telephony issues about which we should be aware?

Advantages

Top 10 reasons for outsourcing:

1. Reduce and control operating costs
2. Improve company focus
3. Gain access to world-class capabilities
4. Free internal resources for other purposes
5. Resources are not available internally
6. Accelerate reengineering benefits
7. Function difficult to manage/out of control
8. Make capital funds available
9. Share risks
10. Cash infusion

Accelerate reengineering benefits
Reengineering aims for dramatic improvements in critical measures of performance such as cost, quality, service and speed. But the need to increase efficiency can come into direct conflict with the need to invest in core business. As non-core internal functions are continually put on the back burner, systems become less efficient and less productive. By outsourcing a non-core function to a world class provider, the organization can begin to see the benefits of reengineering.

Access to world class capabilities
World class providers make extensive investments in technology, methodologies, and people. They gain expertise by working with many clients facing similar challenges. This combination of specialization and expertise gives customers a competitive advantage and helps them avoid the cost of chasing technology and training. In addition, there are better career opportunities for personnel who transition to the outsourcing provider.

Cash infusion
Outsourcing often involves the transfer of assets from the customer to the provider. Equipment, facilities, vehicles and licenses used in the current operations have value and are sold to the vendor. The vendor then uses these assets to provide services back to the client. Depending on the value of the assets involved, this sale may result in a significant cash payment to the customer.

When these assets are sold to the vendor, they are typically sold at book value. The book value can be higher than the market value. In these cases, the difference between the two actually represents a loan from the vendor to the client which is repaid in the price of the services over the life of the contract.

Free resources for other purposes

Every organization has limits on the resources available to it. Outsourcing permits an organization to redirect its resources, most often people resources, from non core activities toward activities which serve the customer. The organization can redirect these people or at least the staff slots they represent onto greater value adding activities. People whose energies are currently focused internally can now be focused externally -- on the customer.

Function difficult to manage or out of control

Outsourcing is certainly one option for addressing this problem. It is critical to remember that outsourcing doesn't mean abdication of management responsibility nor does it work well as a knee jerk reaction by a company in trouble.

When a function is viewed as difficult to manage or out of control, the organization needs to examine the underlying causes. If the requirements expectations or needed resources are not clearly understood, then outsourcing won't improve the situation; it may in fact exacerbate it. If the organization doesn't understand its own requirements, it won't be able to communicate them to an outside provider.

Improve organizational focus

Outsourcing lets the organization focus on its core business by having operational functions assumed by an outside expert. Freed from devoting energy to areas that are not in its expertise, the company can focus its resources on meeting its customers' needs.

Make capital funds available

There is tremendous competition within most organizations for capital funds. Deciding where to invest these funds is one of the most important decisions that senior management makes. It is often hard to justify non-core capital investments when areas more directly related to producing a product or providing a service compete for the same money.

Outsourcing can reduce the need to invest capital funds in non-core business functions. Instead of acquiring the resources through capital expenditures, they are contracted for on an "as used" operational expense basis. Outsourcing can also improve certain financial measurements of the firm by eliminating the need to show return on equity from capital investments in non core areas.

Reduce operating costs

Organizations that try to do everything themselves may incur vastly higher research, development, marketing and deployment expenses, all of which are passed on to the customer. An outside provider's lower cost structure, which may be the result of a greater economy of scale or other advantage based on specialization, reduces a company's operating costs and increases its competitive advantage.

Reduce risk

Tremendous risks are associated with the investments an organization makes. Markets, competition, government regulations, financial conditions and technologies all change extremely quickly. Keeping up with these changes, especially those in which the next generation requires a significant investment, is very risky.

Outsourcing providers make investments on behalf of many clients, not just one. Shared investment spreads risk, and significantly reduces the risk born by a single company.

Resources not available internally
Companies outsource because they do not have access to the required resources within the company. Outsourcing is a viable alternative to building the needed capability from the ground. New organizations, spin-offs, or companies expanding into new geography or new technology should consider the benefits of outsourcing from the very start.

Here are some of the reasons why outsourcing is a better choice for outsourcing the management of a datacenter, rather than maintaining IT internally:

Improved Cost Management Controls
IT costs become more visible as all billable hours must be accounted for Outsourced services are utilized as needed, and organizations pay only for what services are actually used. An outsourced IT department can reduce costs by utilizing its extensive knowledge base of various IT specialists, as opposed to an organization maintaining a comprehensive in-house staff.

Improved Service Quality
Outsourced companies make performance reports and measurements available to their clients. Communications between business functions improve at all levels and ensure that IT resources are not being misused. Outsourced companies can provide 24x7 support at a fraction of the cost. Outsource staff tend to drive planning and budgeting improvements and Service Level Agreements (SLAs) can be established.

Staffing
Staffing levels can be adapted quickly to client requirements, thereby avoiding gaps due to attrition, business growth or economic downturns. IT consultants are fully trained on the latest technologies. Retention of technically qualified in-house personnel is more difficult when external job opportunities are rapidly changing.
Continuous IT support coverage without having to rely on only one or two key people

Equipment and Software/Tools
Outsourced services use established standards for equipment and software requirements, saving time and money. Outsourced services use approved lists of reliable vendors, which improves the quality of goods and services received.

The State of WA can focus on core business competencies
Client IT management can concentrate on core competencies, and improving services for state of WA citizens, while leaving technology infrastructure management to IT professionals

Disadvantages of outsourcing IT

Loss of Managerial Control:
Whether you sign a contract to have another company perform the function of an entire department or single task, you are turning the management and control of that function over to another company. True, you will have a contract, but the managerial control will belong to another company. Your outsourcing company will not be driven by the same standards and mission that drives your company. They will be driven to make a profit from the services that they are providing to you and other businesses like yours.

Potential Hidden Costs: You will sign a contract with the outsourcing company that will cover the details of the service that they will be providing. Anything not covered in the contract will be the basis for you to pay additional charges. Additionally, you will experience legal fees to retain a lawyer to review the contacts you will sign. Remember, this is the outsourcing company's business. They have done this before and they are the ones that write the contract. Therefore, you will be at a disadvantage when negotiations start.

Potential Threat to Security and Confidentiality:
The life-blood of any business is the information that keeps it running. If you have payroll, medical records or any other confidential information that will be transmitted to the outsourcing company, there is a risk that the confidentiality may be compromised. If the outsourced function involves sharing proprietary company data or knowledge (e.g. product drawings, formulas, etc.) this must be taken into account. Evaluate the outsourcing company carefully to make sure your data is protected and the contract has a penalty clause if an incident occurs.

Potential quality Problems:
The outsourcing company will be motivated by profit. Since the contract will fix the price, the only way for them to increase profit will be to decrease expenses. As long as they meet the conditions of the contract, you will pay. In addition, you will lose the ability to rapidly respond to changes in the business environment. The contract will be very specific and you will pay extra for changes.

Tied financial well being to another company:
Since you will be turning over part of the operations of your business to another company, you will now be tied to the financial well-being of that company. It wouldn't be the first time that an outsourcing company could go bankrupt and leave you holding-the-bag.

Possibly Bad Publicity and Ill Will:
The word "outsourcing" brings to mind different things to different people. If you live in a community that has an outsourcing company and they employ your friends and neighbors, outsourcing is good. If your friends and neighbors lost their jobs because they were shipped across the state, across the country or across the world, outsourcing will bring bad publicity. If you outsource part of your operations, morale may suffer in the remaining work force.

Telephony:
Virtualization has emerged as a leading trend in the data center. Organizations are embracing the virtualized data center for a variety of reasons, including total cost of ownership (TCO), consolidation, operational efficiency and flexibility, disaster recovery and business continuity, security, and reduced carbon footprint.

Deploying telephony services on a virtualized platform delivers the following advantages:

- Allows customers to extend the benefits of their virtualization and storage data center investments to telephony services.
- Maintains the same predictable scalability on a physical server.
- You can also run a mix of virtual and non-virtual machines.

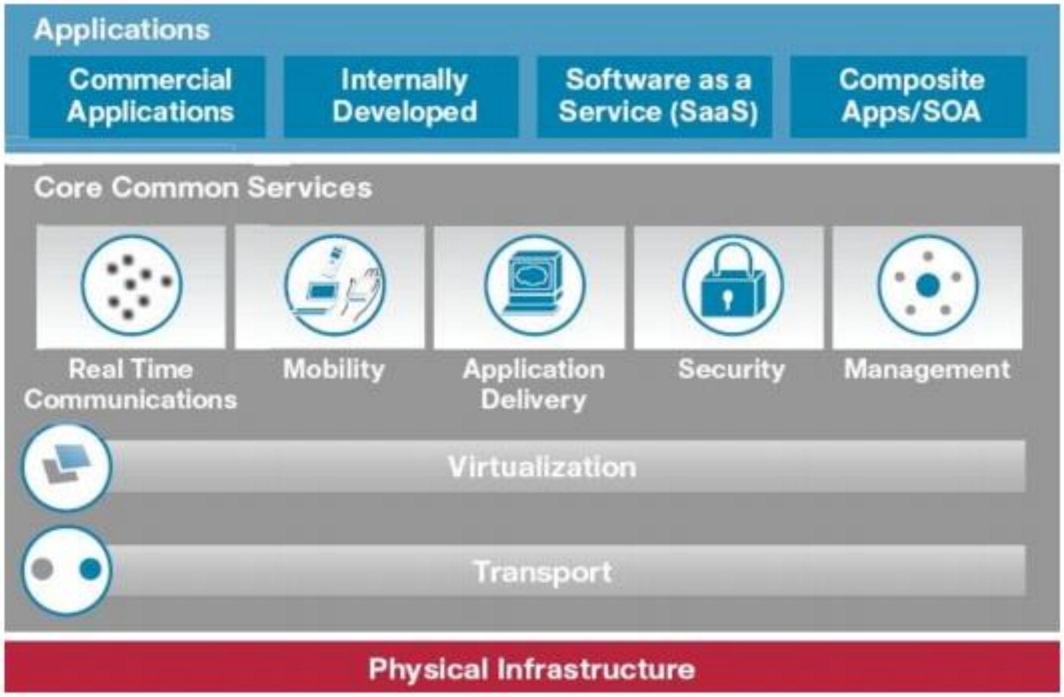

Core Common Services are centered on two principles: application-focus and reusability. This layer comprises seven major core service groups, which deliver consistent and robust capabilities throughout the converged network:

- Real-Time Communication Services offer session and media management capabilities, contact center services, as well as presence functions.
- Mobility Services provide location information, as well as device-independent functionality.
- Application Delivery Services use application awareness to optimize performance.
- Security Services help protect the infrastructure, data, and application layers from constantly evolving threats, and also offer access-control and identity functions.
- Management Services offer configuration and reporting capabilities.
- Virtualization Services deliver abstraction between physical and functional elements in the infrastructure, allowing for more flexible and reliable service operations and management.
- Transport Services help with resource allocation and deliver on the overall QoS requirements of the application, as well as routing and topology functions.

Deploying IP telephony services on a virtualized platform delivers the following advantages:

- Allows customers to extend the benefits of their virtualization and storage data center investments to telephony services
- Extension of disaster recovery capabilities inherent to virtual infrastructures for telephony core services
- Maintains the same predictable scalability on a physical server

You can also run a mix of virtual and non-virtual machines

Question 2: How would this deployment strategy address unique aspects of the centralized state data center such as security concerns and disaster recovery? For background on state standards, see the Information Services Board: http://isb.wa.gov/policies/Default.aspx

The virtual hosting and recovery approach INX proposes as the shared services turnkey strategy inherently addresses security and disaster recovery aspects using the most advanced technologies in the market today.

Security Aspects

System Management

The state of WA Shared Services requirements states a need for centralized management capability across security gateways of various types and form-factors. Conceptually, the management solution is intended to provide a "single pane of glass" for configuring and operating most, if not all, network security enforcement points. DIS currently provides customers segregated security domains or "contexts" to streamline administration and allow for agency interaction with their associated security policies. The virtual hosting service (VHS) management system should maintain and expand on this capability to provide "virtual" firewalls within scalable security appliances and allow for the extension of granular role based access to customers.

DIS anticipates the need for hypervisor-level firewalling to control inter-virtual machine traffic. The INX VHS solution would allow for unified management of policy enforcement at the customer perimeter, SDC network border, SDC services edge, and at the hypervisor.

System Visibility

The OB2 Move Project Charter calls out the following high-level requirements and INX assumes similar requirements are relevant to the shared services cloud:

- SDC Event Correlation/Root Cause Analysis process definition and tools are required to conduct Problem Management.
- Umbrella Tools which tie system alerts together are required to conduct Problem Management.

The VHS security architecture will include a Security Information and Event Management System (SIEM). The SIEM solution will receive security and change control events from all deployed SDC core network and security components. The containerized architecture must contemplate the phased addition of an extended event scope including: physical security, network, server and application events. The solution needs to scale to meet the anticipated increase in event scope without significant impact to the core VHS SIEM infrastructure.

Flexibility and Scalability

Current SDC network reference architecture sets 10 Gigabits per second as a standard speed for converged networking across the SDC core. Uplink speeds to the Internet are currently 500 Megabits per second. Individual security components should support current loads and anticipated growth without the need to replace appliances prematurely. Security devices would support a "scale-out" model to accommodate unanticipated demands and new service offerings. The modular architecture can extend to multiple disparate environments to apply a common portfolio of security services.

Compliance

The VHS security architecture must contemplate agency use cases that include multiple regulatory requirements and audit processes. In addition to State of Washington ISB Policy, current regulatory requirements include: HIPAA/HITECH, PCI-DSS, IRS 1075/NIST 800-53, and FBI requirements. The architecture will be capable of supporting progressively more stringent security requirements without constraining customers that are only subject to ISB Policy.

Network Instances

Internet MPLS VPN

The Internet VPN routing and forwarding (VRF) will be provisioned as a dedicated separate routing table on the core devices. It will provide a connection point for key Internet facing services.

The creation of this network, segregated logically from all other traffic by MPLS technology, offers the ability to isolate secure services from unsecure services, and allows 'risky' traffic to be kept off devices in the data center.

Transport VRF

The transport VRF provides a semi-secure network for interconnecting state services. No systems should be directly connected and it will be protected against Internet-borne attacks by a layered defense.

Data Center

A variety of networks will exist inside the SDC and VHS, including containers for agencies, shared services, and legacy DIS services. These will connect, via the data center security layer, to the Transport VRF.

Containers will also be extended to the appropriate agency network. This is particularly important where the agency has services and servers located in the Data Center; these will, effectively, be part of the agency network via this container extension operation.

Management Networks

Management will be via Out of Band management networks. These will run across the same infrastructure as the other networks, but will be in separate VRFs and terminated into a secured zone. All management, logging, and authentication activities will occur across this network.

Customer Networks

Customer/Agency networks will be delivered as their own VRF instances across the infrastructure. They will remain separated and can be delivered to the data center, allowing containers of customer systems to effectively be internal to the customer organization.

Internet Facing Elements

The VHS architecture provides for a layered perimeter. The Internet edge of the VHS network provides minimal filtering; it effectively provides a connection to an internal connection point for Internet services. As such, 'edge' security is minimized. This allows for further development of identity based services as the traditional security border becomes less and less effective.

Internet Perimeter Security

The Internet Perimeter security is designed to protect the infrastructure for coarse grained attacks and provides a limited level of security to the Internet VRF. It will provide stateful packet filtering and limited application awareness. A purpose built security appliance will provide modularity at the internet edge and offer the Security team a stronger feature set for future business needs.

Internet DMZ Services

The Internet DMZ services are a continuation of the existing 'PGN' services which cannot be proxied through the Secure Access Washington infrastructure and must be presented 'raw' onto the Internet. The DMZ hosts will be built on dedicated physically segmented server hardware and will connect directly to the Internet VRF via a stateful packet filter with application awareness.

Gateway Services

Gateway services provide an interface between the unsecure Internet VRF and the semi-secure Transport VRF and will provides the 'hub' of the network, interconnecting all the services inside a perimeter.

These gateway services will provide high speed transit between the Internet and critical services that connected via the transport VRF.

Proxy Services

The proxy services currently provided by DIS ESS will be provisioned into a dedicated network and will be firewalled from both the Internet VRF and the Transport VRF. It is not expected that any component of this area will participate in any routing table; data will be handed from the Internet VRF to this domain, and then passed from the secure end of the proxy to the Transport VRF.

Remote Access

Remote access systems will be directly connected to the Internet VRF and filtered on the Internet perimeter. These systems will be sufficiently hardened to exist in this environment without additional firewall proliferation.

Backend access from these systems will be mapped either directly to the customer VRF or to the transport VRF. Security on the VPN devices will be enforced by identity based technologies that are tightly integrated to the authentication.

IGN connectivity
The IGN will connect to both the Transport and Internet VRF's through dedicated firewalls, managed by DIS. All data will be subject to application and vulnerability inspection. A single security appliance will be provisioned to enforce this and will not participate in routing.

Transport VRF Edge
The transport VRF edge is the primary security for State users heading to the Internet. Inbound traffic will be secured by one of other gateways, be it the proxy infrastructure or remote access. By defining this as an egress-only control point, the number and type of controls can be minimized.

Secondarily, these controls are designed to enforce policy upon the user community, ensuring that outbound connections are appropriate. This includes URL filtering as well as application control.

Customer Edge
The customer edge provides for the connection of end-users to the Transport VRF. These customers will be secured (according to their security policy) from traffic originating in the Transport VRF and will be able to control traffic exiting TO this VRF. Customer traffic will be brought through the core via MPLS and terminated in the Security Layer. Any active routes exchanged with the customer will be filtered at the network level to prevent accidental or deliberate black-holing of other traffic on the Transport network.

Data Center Services Elements
All of the above services will reside in the SDC. The services and compute aspects of the data center (here rolled up simply under 'services'), however, need additional consideration.

There are two classes of 'container' in the architecture. The first is an A'la carte customer, hosting their services inside the SDC. These customers will not take advantage of the Data Center Service Security but, rather, will have their containers mapped to their networks via the infrastructure. The security of those containers is the responsibility of the customer.

The second class is DIS managed and maintained systems (H2-M2), such as the shared services applications. These containers will require security prior to being connected to the Transport VRF for access by clients. It is these areas that this architecture focuses upon.

Business continuity is a strategic initiative in which the business owner's determine and outline requirements to ensure successful recover in the event of a disaster. This included determining what personnel and resources are necessary for the IT infrastructure to survive should a disaster occur. The outcome is a business continuity plan.

Disaster recovery is a tactical IT initiative for the IT department to meet business continuity requirements. As with business continuity, the goal is to determine who and what is necessary for the IT infrastructure to survive a disaster. The outcome is a disaster recovery plan.

INX recommends we leverage existing ISB DR-BC policies, standards, and departmental DR procedures, to drive RTO and RPO requirements for the shared services infrastructure. In the case where departmental DR plans are inadequate or non-existent it is of critical importance these plans be defined. The recommended approach to address inadequate DR plans is to contain the scope of these efforts to business critical and line of business departmental applications that are near term virtualization targets.

Redundancy:
Component redundancy for each of the infrastructure and hardware components is just as important in the virtual infrastructure as it is in a physical infrastructure. Multipathing for storage, redundant network links, redundant hardware devices and an N+1 server architecture are all methods for providing redundancy in the Shared Services environment.

Backup and Recovery Strategies
Several data-protection strategies are available including layer 2 extension between the use of guest OS-based software, VMware Consolidated Backup, VMware Data Recovery, VMware snapshots, storage snapshots, and storage replication. VMware Site Recovery Manager is a business continuity and disaster recovery workflow product.

Staffing and facilities do not depend on whether a solution is physical or virtual other than the need for training and the facilities components to support the DR requirements.

Guest OS Backup: Guest OS backups operate identically to physical machine backups. A backup agent is installed within the virtual machines and is registered with a backup server. Recovery, just as a physical machine, can be done in either of two ways:

- Bare-metal restoration, if supported by the server vendor, is one option. This is optimal for guest OS-based backups since it does not require an OS installation.
- Installation of an OS and backup agent prior to performing restoration is the other option

VMware Consolidated Backup/vStorage API: VMware ?Consolidated Backup enables offloaded and impact-free backup for virtual machines on an ESX host by allowing traditional file based backup software to leverage virtual machine snapshot technology and efficient SAN-based data transfer. Consolidated Backup is a backup enabler that allows for LAN-free backups handled by a SAN-connected proxy server. The Consolidated Backup proxy server handles backup processing, which frees resources on the virtual machines and the ESX host.

Consolidated Backup performs two types of backup:

1. File-level backup
2. Full virtual machine backup

File level backups are similar to traditional backup methods in that full and incremental backups can be completed for Microsoft Windows OS. Full virtual machine backups are similar to bare-metal restore backups and include all of the virtual disks, configuration files, NVRAM file, and log files for a virtual machine.

Virtual Machine Snapshots: Virtual machine snapshots take a point–in-time copy of a virtual machine, including disk, RAM, and the virtual machine configuration. In the event of a guest OS failure and administrator can revert to one of the previously created snapshots very quickly.

An important consideration when utilizing snapshots is that Consolidated Backup prefers that a virtual machine have no snapshots. It is also important to keep in mind that a snapshot requires disk space to store a copy of the virtual disks and other information from the snapshot. This is a critical Shared Services infrastructure design consideration that must be handled with rigor to assure successful application continuity.

Planning and designing the storage and replication for a VMware vSphere infrastructure is closely related to larger disaster recovery planning efforts. Several replication strategies provide different levels of service and associated costs. Fort most organizations, a multi-tiered replication strategy provides the most cost-effective VMware vSphere infrastructure replication solution. Organizations typically classify approximately 80% of their servers and data with a 24-hr RTO and the remaining systems with stricter RTOs, ranging from several hours down to continuous data protection (CDP) with synchronous wide area replication.

As with any disaster recovery plan, the RPO and RTO are constrained by budget. VMware vSphere enables extremely short RTOs, because servers can be restored and placed in service very quickly. Indeed, it is often possible to have virtual machines up and running within minutes following a disaster declaration. RPO, on the other hand, is constrained by bandwidth. Replicating a large amount of changed data over a small data pipe increases the RPO. This makes the selection of a storage replication strategy critical to a successful VMware vSphere infrastructure DR plan.

Storage Replication Strategies:

Several replication strategies for the Shared Services virtual infrastructures align nicely with the backup and recovery system tiers and technologies.

The 80% of servers and data that are backed up daily with a traditional 24-hour RPO can be replicated through a variety of means from the main site to the recovery site. This type of RPO is often best met by replicating backups. After the nightly backups are completed (file or vmdk-based), the data needs to be replicated to the recovery site before the next backup window begins. When replicating a large volume of data, it can become a challenge to maintain eve a 24-hour RPO at a level of bandwidth that does not exceed the budget.

The 20% of remaining systems can have much stricter RPOs, some even requiring continuous data protection and transactional replication. In these cases, replicating nightly backups is insufficient. Either asynchronous or synchronous replication between storage mediums is required.

Several strategies exist for enabling storage replication, depending upon RPO and budget. The primary differentiation between these replication strategies is where they sit in the storage stack.

Embedding storage intelligence directly into the storage fabric creates the highest performance and transparency. Virtual SANs, LUN virtualization, remapping, and replication can all be enabled to the Shared Services infrastructure. Synchronous remote replication of some or all of the storage traffic is accomplished by using a variety of storage replication protocols.

When the higher-performing fabric-based replication solution is not practical or economic, the recommended solution can be man-in-the-middle replication server or appliance to provide asynchronous in-band or side-band replication of LUNs, and virtual machine exports. These solutions include software or appliance-based storage virtualization or replication systems (Veeam, esXpress). They can also be used for asynchronous mirroring of operating system or vmdk-based backups on disk, virtual tape library, RDM LUNs or NAS storage. This replication strategy is still transparent to the VMware infrastructure but not to the storage infrastructure.

The man-in-the-middle replication scheme is most effective when combined with data deduplication solutions (NetApp, Avamar, Data Domain). Data deduplication can dramatically reduce the amount of replicated data, particularly when replicating backups that involve a great deal of duplicate data. Adding data bandwidth compression further enables rapid replication with less bandwidth required.

Question 3: What scale of operation would be required for this alternative to be attractive to potential bidders on any future RFP? Is there a critical mass below which you would not be interested in bidding on a proposal for this approach? What metrics would capture that decision?

After many years of learning lessons in large datacenters throughout the world, INX engineers and architects have learned that scaling in smaller server groups ensures project success. INX recommends virtualizing incrementally as state agencies adopt server consolidation strategies. An incremental approach helps mitigate investment risk and allows for DIS processes to mature.

From a hardware stability perspective, acquiring servers in batches no larger than 300 physical servers at a time, avoids inconsistency in manufacturer firmware and part models, speeding up racking times and making firmware and software update efforts transparent, manageable and repeatable. In smaller provisioning batches, DIS can better track server provisioning timelines and inventory stacks versus resource allocation in the infrastructure for capacity planning purposes.

As a global best practice, consolidated virtual server infrastructures become unmanageable after the 300 virtual machine inflection point. Provisioning, governance, enterprise change management, and capacity planning practices combined with rich monitoring and troubleshooting tools become critical in ensuring virtual infrastructure initiative success as the architecture grows. Best practice containerized cloud infrastructures will scale, but the adoption of ITIL-based process methodologies must also scale accordingly. Cloud computing represents a dramatic acceleration of every stage of the service delivery process.

Given an internal cloud infrastructure— typically a high-end computational platform in combination with virtual server management, system monitoring and provisioning tools—the time required to create new services can fall from multiple weeks to less than a single business day with a lightweight provisioning governance model.
As the containerized virtual infrastructure scales, the cost per virtual machine can drop significantly.

Virtual Machine Growth Cost Curve

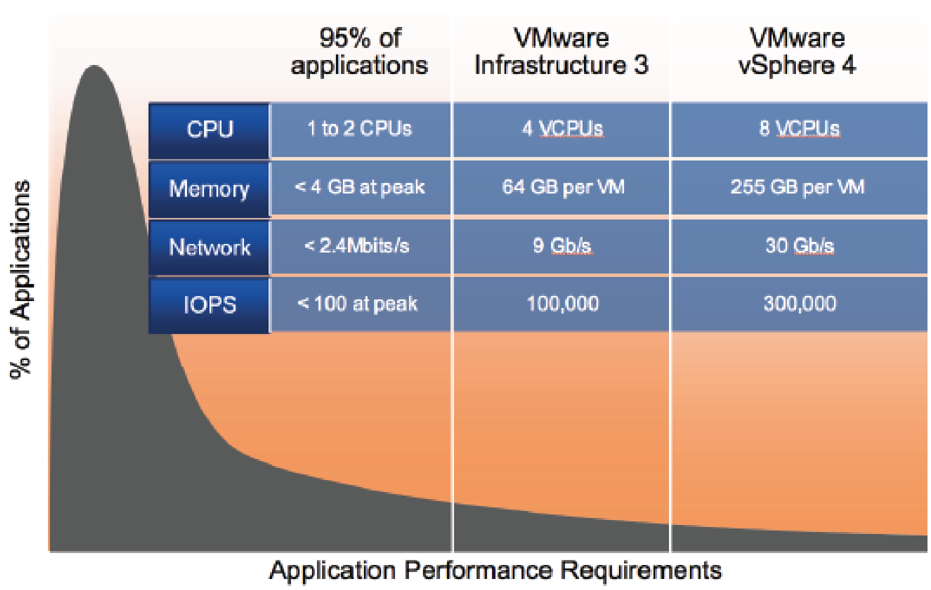

Virtual machine CPU count has doubled from four to eight with vSphere 4. Host CPU core maximum has similarly doubled from 32 to 64. Maximum host memory has quadrupled, from 256 GB to a full terabyte; the amount of memory that can be allocated to any single virtual machine has also quadrupled, from 64 GB to 255 GB. VSphere now supports up to 320 virtual machines and 512 vCPUs.

The increased processing power of more cores and additional memory capacity allow a shared services organization to virtualize larger business-critical workloads without impacting application performance. The more applications that can be supported in a virtualized environment, the greater the reduction in cost, risk, and downtime. These increases in capacity allow architects to design virtual hosts conservatively for performance while leaving headroom to scale for rapid provisioning of new application services.

Question 4: What intellectual property issues, if any, would limit state's ability to pursue this approach? Any proposed solutions must be based on the appropriate technical solution and equipment as opposed to the lead vendor's default proprietary solution.

Intellectual Property:

No matter what the technology or service being outsourced, the IP rights of all parties to the transaction – and perhaps of third parties – will likely be involved in one form or another. Out of that involvement comes the need during contract negotiations for both the customer (DIS) and the outsourcer (INX) to focus on and control the intellectual property itself. Also, during the subsequent outsourcing relationship, issues involving intellectual property will likely arise that will require ongoing resolution.

Outsourcing allows an entity to transfer portions of its internal business functions, such as IT infrastructure, personnel, business processes and applications, to another resource outside of the organization ('outsourcer'). Offshore outsourcing ('offshoring'), a hotly debated and political issue, may involve the transfer of intellectual property to foreign countries and the elimination or transfer of current US employees to overseas locations, where salaries may be markedly lower and intellectual property may not be as protectable as it is in the United States. (Note: Offshore outsourcing is not recommended.)

Intellectual property transferred to the outsourcer (INX) may include software, data, business and technology processes, trade secrets and other confidential information, inventions, patents, trademarks, copyrights and other works of authorship, and know-how. Some of this intellectual property may belong to third parties and be licensed to the customer.

IP audits have been used over the years for various purposes. Results of the IP audit would be used to make sure that the client's intellectual property is protected properly, to maximize return on investment, to make sure that the client is not violating third-party rights and to assert appropriate IP rights against those who would violate such rights.

Once the IP audit is completed, the results can be used during the negotiations of the outsourcing agreements to schedule the intellectual property and deal with the usual and unusual IP-related issues that may arise. In addition, the IP audit results can be used to make certain that relevant intellectual property is treated properly and protected.

In addition to the customer IP audit, it is also important for the customer to get a clear understanding of any outsourcer intellectual property that will be used during the course of the outsourcing relationship, and to learn how such outsourcer intellectual property might impact on, interact with or otherwise affect use of the customer intellectual property.

INX does not offer a proposed solution and expects to collaboratively solve this issue if selected to continue with DIS.

Section 3 -- Responses to Questions regarding the Transition

Briefly address each of the four (4) questions in the Transition section of the RFI description.

Question 5: Given the current decentralized allocation of resources, what role would the transitional data centers in agencies and the statewide data center in Office Building 2 play in your proposed alternative?

The current Washington state computing environment is largely comprised of separate computing environments for each agency. In many cases, these computing environments are comprised of similar hardware and software, tools and a complete compliment of people necessary to architect, acquire, construct and operate these systems. The primary differentiation of these environments is in the applications and not in the underlying infrastructure that supports these applications. Despite the inter-agency similarities, there is little exchange of expertise, tools, or innovation and little leverage of the investment. By abstracting the applications from the underlying infrastructure and operating systems, The Shared Services initiative team at the state of Washington recently completed reference architecture for a scalable multi-tenant cloud infrastructure to resolve these inefficiencies. INX proposes to build upon those architectures to achieve an end state that meets or exceeds DIS expectations.

Nearly every agency in the state of WA has some level of virtualization in place today. INX proposes to leverage that progress and rapidly accelerate the adoption of virtualization technologies in each of the agencies in preparation for transition to the SDC. Preliminary ROI figures for a recently proposed project for the Dept of Ecology suggest a 5 year $1.8M ROI by virtualizing 115 servers and associated storage. Regardless if DIS undertakes and outsourced model, INX recommends aggressively pursuing this level of virtualization in all agencies for the financial benefits that are needed at the state today.

The next phase in optimizing the datacenter is consolidating and virtualizing all possible servers in preparation for the move into the new State Data Center. INX proposes a multistep approach to enable the agencies to prepare for the Shared Services infrastructure and reduce capital and operating cost in the near term:

Solution Overview

At a high level, the INX solution follows an Analyze-Design-Plan-Execute methodology. INX will run a virtualization capacity planner to gather facts and data about each agencies physical and virtual infrastructure to determine consolidation scenarios for this virtualization project. Additionally, INX will run storage scanning tools to predict the amount of storage recovery upon completion of data deduplication. The outputs of these reports serve as an input to the Design and Plan phase.

Key solution considerations include:

- Look for opportunities to leverage existing infrastructure investments
- Provide recommendations for an optimized virtual storage infrastructure
- Enable a foundation for a virtual replication DR strategy
- Look for ways to increase agency storage efficiency by leveraging deduplication

Cloud Computing Challenges Today

Disregarding the SaaS cloud providers such as Salesforce.com which effectively enable selective Internet based applications; there are two primary models for cloud computing today: trusted providers and public providers. Trusted cloud providers are either internal or external data centers that provide resources on a rental or as-needed basis along with an engineered data transport architecture (i.e. fiber, T-1, MPLS). Public providers, such as Amazon's EC2, utilize a purely Internet-based access model of shared resources via proprietary APIs.

As **Figure 1** show, the risk increases significantly when accessing infrastructure from public versus trusted providers. The disadvantage of public providers is that the proprietary APIs lock customers into the providers' view of control. The Internet-based access entails packet transport in unpredictable ways, leaving the potential for large amounts of latency. The combination of public provider proprietary API/environment along with Internet connectivity eliminates all customer control over ability to effectively meet SLAs and to comply with security requirements built on defense in depth strategies. Another challenge public cloud providers face is the lack of multi-tenancy containment strategies that isolate client compute, data storage infrastructure, private LAN services from shared cloud interconnect points in the architecture.

Trusted cloud providers allow customers to maintain various levels of control depending upon their agreement with the trusted advisor by continuing to maintain access to, or even ownership of, equipment, software and connectivity. But the trusted provider model is severely hampered by lack of an efficient means to easily access those resources as a seamless part of a customer's primary data center architecture.

Figure1 - Risk Trade-Off between Trusted and Public Cloud Computing Providers

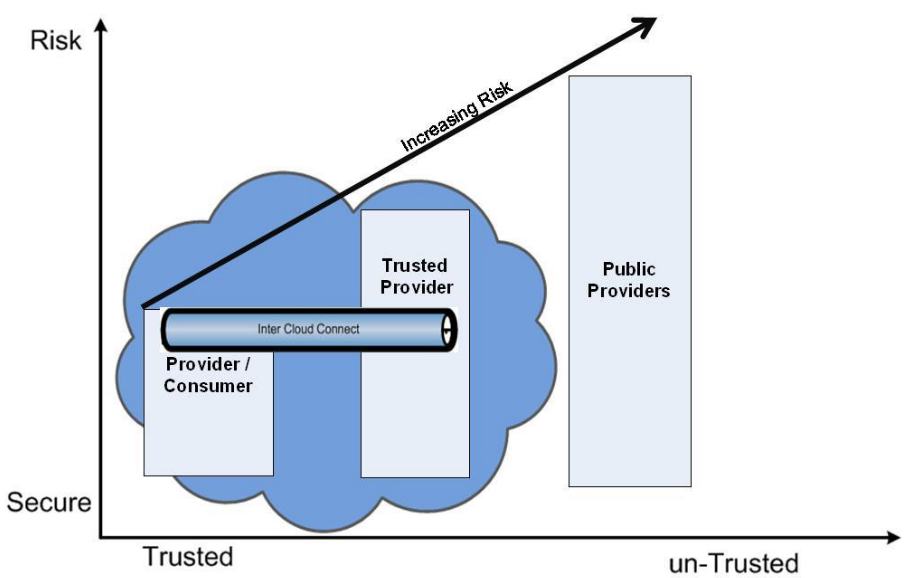

The only way to seamlessly integrate a customer's primary data center with a trusted provider's resources is to utilize virtualization architecture along with virtual machine live migration (vMotion) between data centers. Several manufacturers are beginning to provide long-distance vMotion solutions, but they inevitably take the approach of utilizing Layer 3 (IP routing). The guaranteed delivery of IP is accomplished with checking and rechecking of transport along with retransmitting of packets which adds a lot of latency and inefficiency. The result is that the solution is typically limited to very expensive 10GB bandwidth connectivity, making it out of reach to the majority of potential customers.

Beyond the transport inefficiencies, a Layer 3 approach to Live Migration entails several other challenges as well including:

- Layer 2 procedures are used to manage VM migrations, but IP WANs block Layer 2 protocols.
- Servers must reside on the same IP subnet to support inter data center mobility
- ARP is required to update the network of IP address moves
- LAN & SAN jumbo packets are not supported by a WAN
- Inefficient tunneling techniques (i.e. GRE, L2TP & MPLS) make cloud transport even more inconvenient
- WAN protocols are intolerant to the packet loss or re-ordering inherent in TCP/IP
- SAN protocols must be enhanced to extend over long distance
- WAN traffic may need to be encrypted for compliance

Extend2Cloud™ Alternative – Transport Container Definition

INX's Extend2Cloud™ (E2C) significantly enhances cloud computing by enabling efficient access to trusted cloud providers, whether internal or external data centers. E2C is built upon a hardware based packet processing engine that consolidates the functionality of the OSI Reference Model layers one through four. The packet processing engine is built on a set of Field Programmable Gateway Arrays (FPGA) that can compress data streams, provide forward error correction at the aggregate level of every ten frames sent, secure transmission via AES 256 bit encryption and performs MTU transparency to send jumbo frames or greater thru an Ethernet tunnel. The hardware engines can be configured in a Point-to-Point or Multi-Point configuration. The WAN interfaces of the hardware appliance can also be configured to create layer two Ethernet rings via Provider Backbone Transport Engineering – IEEE 802.1AQay PBB-TE E2C presents storage in geographically disparate data centers as if it were one logical data center irrespective of distance.

As **Figure 2** shows, E2C provides an encapsulation/abstraction layer between the HW/SW technologies of an enterprise data center.

E2C follows the telecommunications carrier model of transport by only utilizing Layer 2. E2C accelerates and guarantees delivery of Ethernet frames in a deterministic fashion by natively extending Layer 2 vLANs. The E2C hardware appliance extends the deterministic nature of Layer 2 across the WAN. WAN solutions could be comprised of dark fiber, TDM (Time Division Multiplexing) over Ethernet, MPLS or even the Internet. The Layer 2 end-to-end transparency enables seamless access of resources between trusted data centers. At VMworld, for example, we demonstrated running a full-motion video VM in one data center while the data store resided in another and the client was on the VMworld showroom floor.

The forward-side error correcting capabilities of E2C combined with today's more efficient transport mechanism enables packet transmittal without loss. It also eliminates the requirement for the TCP slow start handshaking algorithm that slowly ups the bandwidth in TCP/IP. E2C injects a check sum on every 10^{th} frame across the WAN which allows every 10 frames to be reconstructed on the far side without issuing a retransmit. In very simple terms, this makes layer 2 reliable without TCP/IP.

Figure 2 - INX Unified Data Center Abstraction Containers

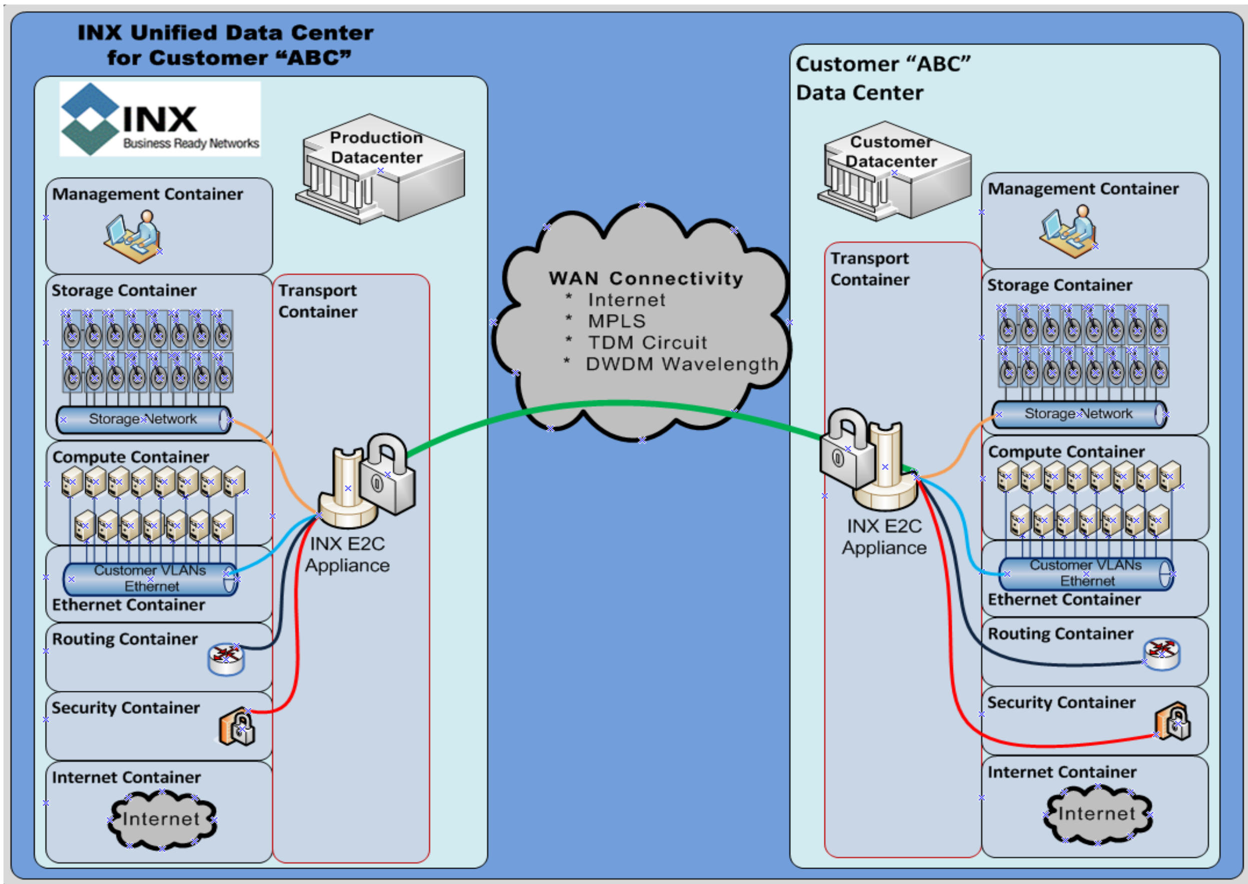

E2C simplifies wide area connectivity by lowering the cost of IP, MPLS and Ethernet WANs and by providing multi-protocol transport for both Ethernet and Fibre Channel. All critical data is encrypted in real time with AES 256-bit encryption using CryptoWIRE, and packet loss and other effects from WAN errors are eliminated. The result is an end-to-end virtual machine migration that reduces protocol latency and round trip delay. Increased throughput is enabled along with VMware High Availability and even Fault Tolerance without increasing bandwidth. **Figure 3** shows the relationship of trusted and untrusted resources through the cloud.

Figure 3: Relationship of Trusted and Untrusted Resource Access

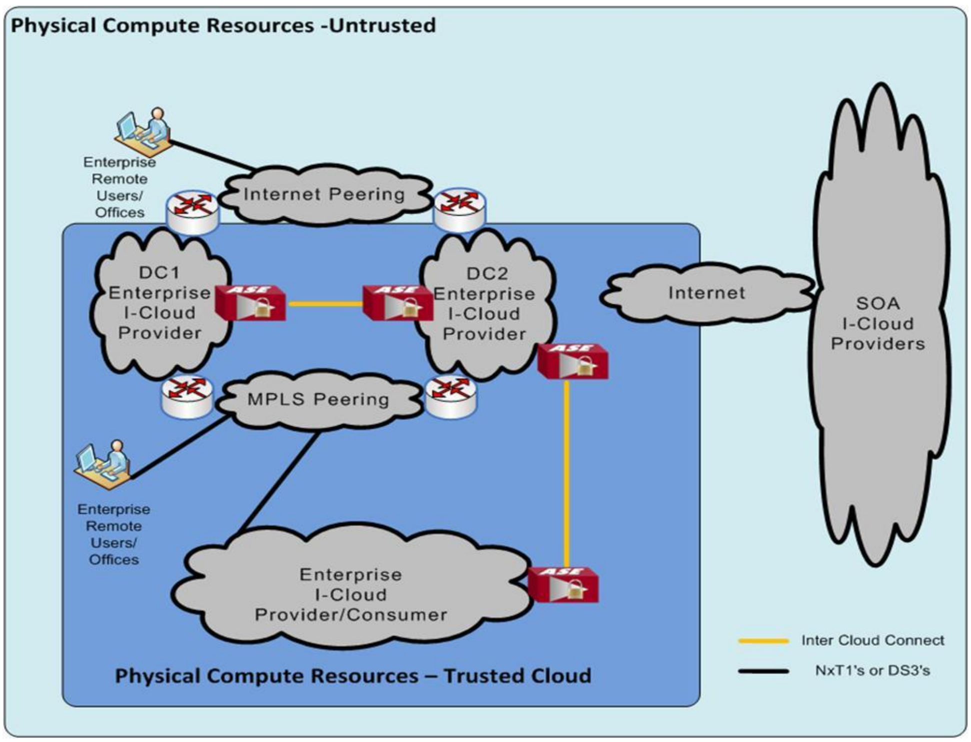

E2C Compression

Effective compression is the final component making the efficient transport of Layer 2 across multi-protocols and long distances possible. WAN accelerators accomplish compression with pattern matching, meaning they are effective only when a bit pattern can be identified. Virtualization, though, utilizes unstructured data in block level storage, meaning that the patterns are for the most part not identifiable. This makes it very difficult to effectively compress byte-level changes in real time.

E2C alternatively uses front-side error correction to prevent requiring re-transmittal of a TCP frame. This eliminates the sliding window algorithm along with the accompanying ramp-up time normally required in order to build to full-speed transmittal. We "Storage vMotioned" a 41GB volume in order to test the increased efficiencies of E2C compression. When decreasing the bandwidth by a factor of 10 (1000Mbps vs. 100Mbps, the transmittal time only increased 162%...from 13 minutes to 34 minutes.

Inter-Cloud Connect / Long-Distance V-Motion

Figure 4 shows the Extend2Cloud™ architectural components demonstrated at VMworld in San Francisco in September, 2009, while **Figure 5** shows the logical components enabling long-distance VMotion to Europe demonstrated as part of VMworld.

Figure 4: Extend2Cloud™ Architectural Schematic from VMworld 2009 – San Francisco

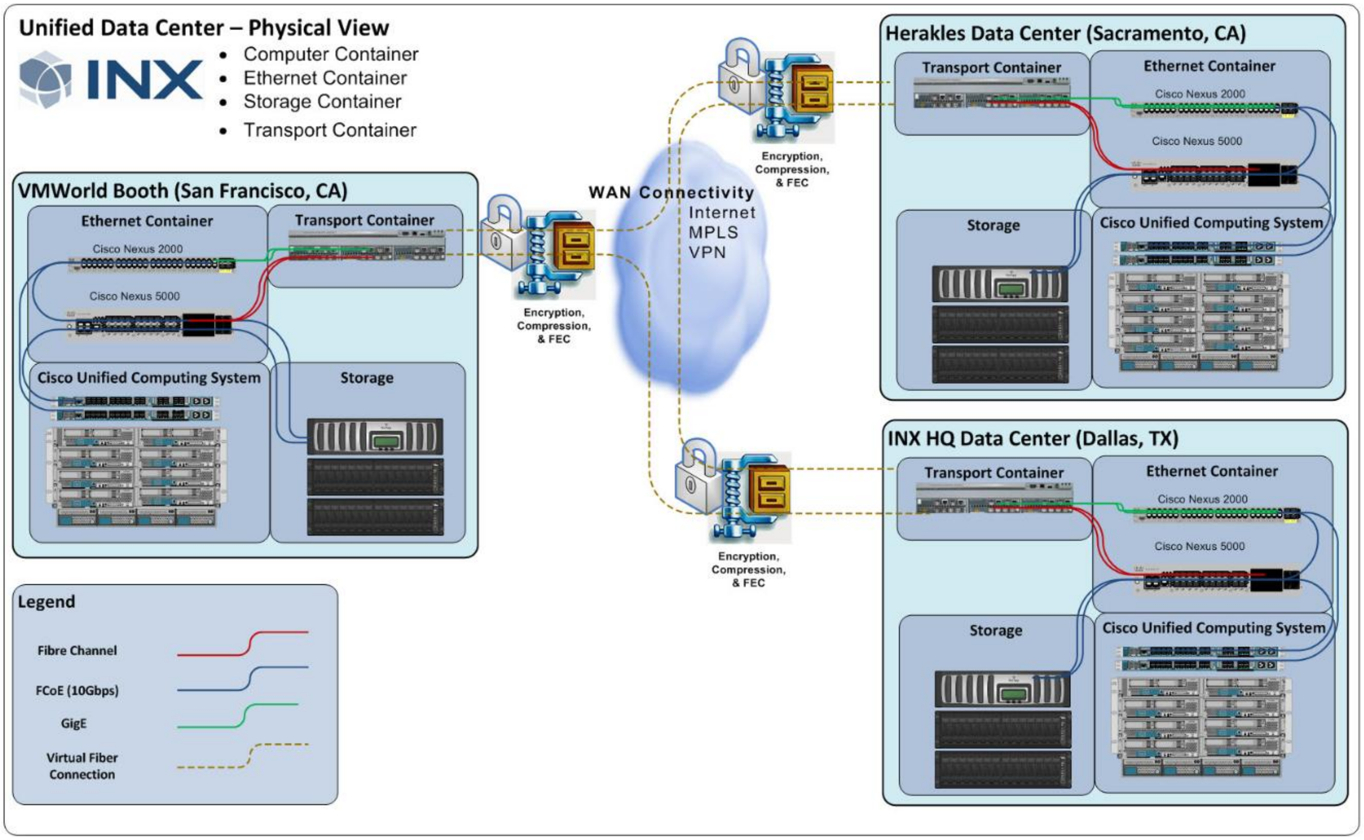

Figure 5: Logical Representation of a Unified Data Center. Demonstration example of E2C concept at VMWorld 2009.

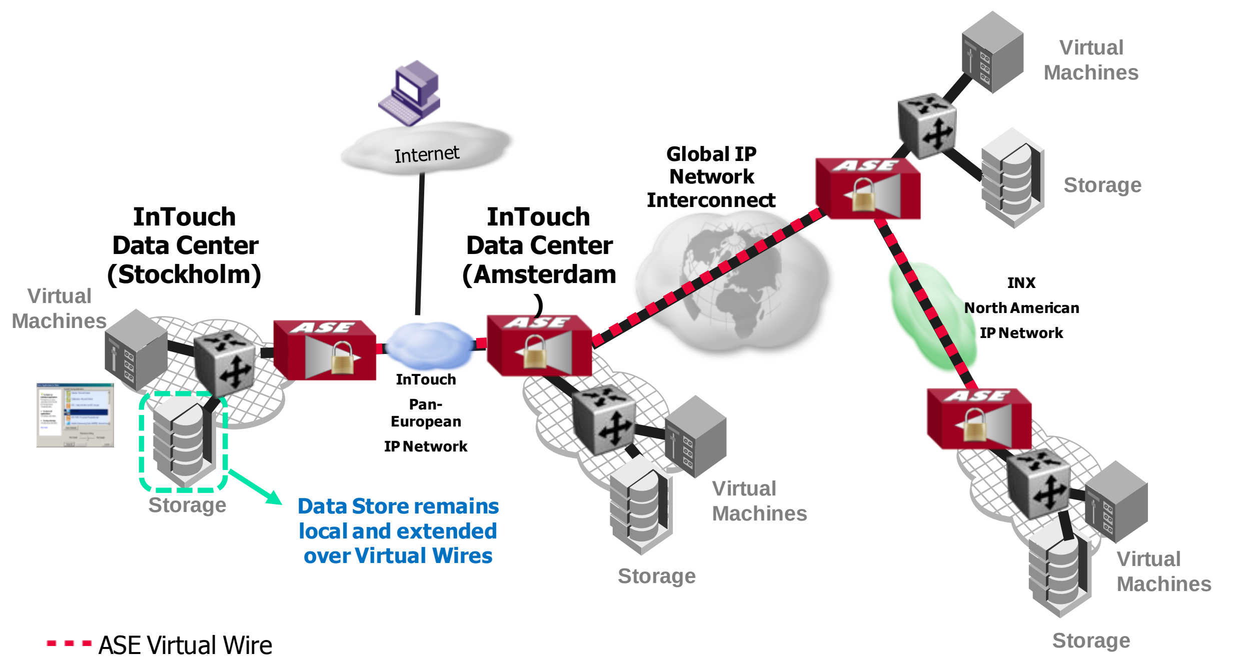

Question 6: What is a potential timeline for the transitioning of decentralized assets to the data center, equipping of the center, start up of the joint asset, initial operation and training of state employees, and transition of the data center to full state ownership and operation? How would the vendor establish the timeline and what elements would be needed to accurately develop a timeline?

INX recommends a strategy of incrementally virtualizing agencies immediately, but do so in a manner that enables virtual machines and application data to migrate over-the-wire to the SDC cloud infrastructure during the OB2 data center move timeframe. Technologies are available today that can enable this approach. The potential timeline below represents a possible sequence of events. Major factors that could impact this timeline include:

- DIS adoption of a rigorous holistic design methodology
- DIS adoption of holistic procurement and contractual approach
- The rate at which WA state agencies virtualize
- OB2 datacenter move schedule delays
- Level of collaboration between DIS and state agencies during the initial virtualization period
- The rate at which state employees develop skill sets for cloud operation and optimization
- Timely network, storage, compute and security product standardization and policy setting
- Timely establishment of a shared service catalog and lightweight governance practices

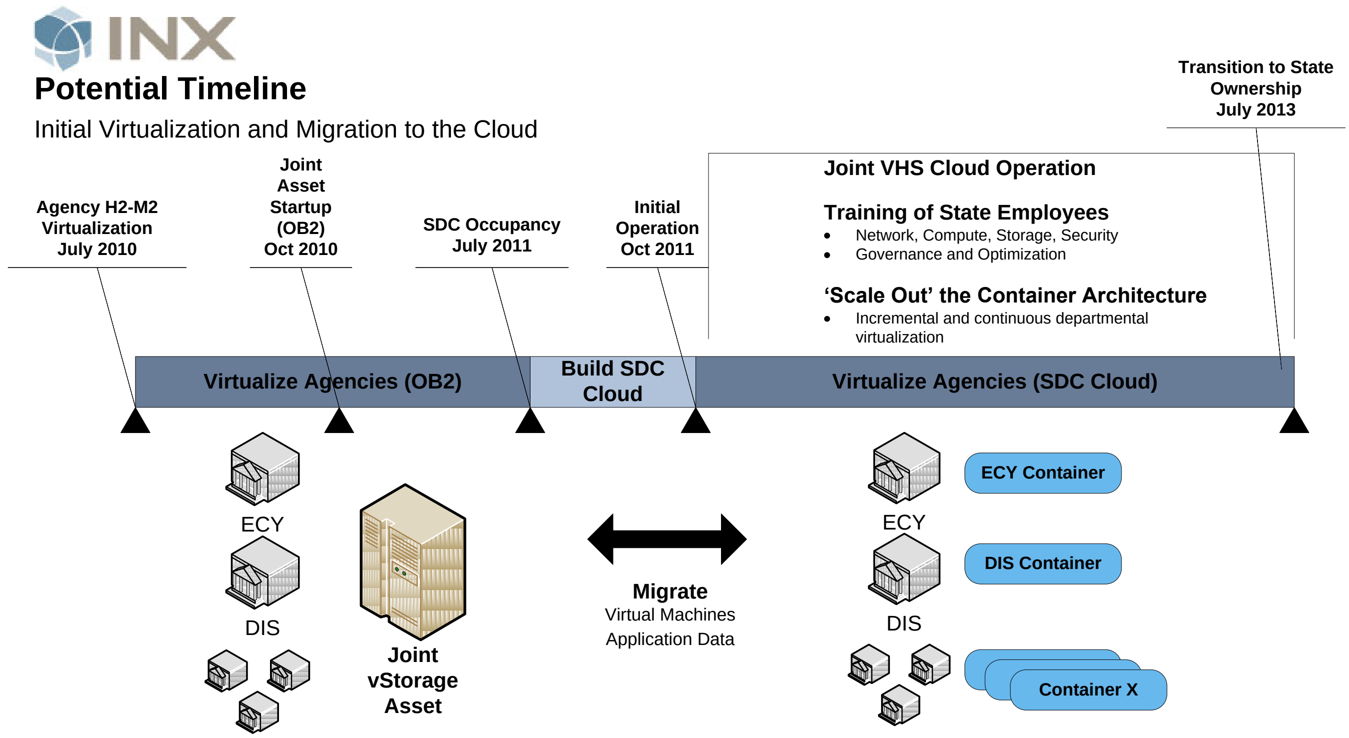

Question 7: How would this approach best align or integrate existing legacy systems and mainframe operations with the proposed private cloud for the public sector? How would a vendor handle the continued coexistence of legacy and modernized solutions in transition?

The set of architectures and technologies collectively referred to as the mainframe has been in existence (and therefore development) for more than forty years. During that time, previously unimagined levels of performance and reliability have been achieved, and the mainframe—written off many times over the years—has continued to be the backbone of corporate computing in many of the world's largest companies. Mainframes represent a major computational capacity within OB2 and other disparate agency datacenters across the state of WA. There are currently three classes of storage in OB2. Those classes are:

1. Mainframe storage supported by EMC's Symmetrix platform,
2. Open system storage currently supported by EMC's CLARiiON platform
3. Content addressable storage (archive storage) supported by EMC's Centera platform.

Prior to the OB2 move into the SDC, the following is planned for the storage build:

- New EMC Symmetrix storage frames will be installed and configured in SDC.
- New open storage frames from an undetermined vendor will be installed and configured in SDC.
- New open storage frames from an undetermined vendor will be installed in OB2 prior to the move to support current disk space needs. The data on this storage frame(s) will need to be migrated and the frame(s) physically moved to SDC.
- New storage area network (SAN) equipment will need to be installed and configured in SDC. The SAN will need to be extended so that the SAN in OB2 and SDC appear as a single SAN allowing data access from any location.

To maximize the return on investment for the state's existing investments INX proposes a design strategy that encompasses all three classes of storage in a holistic approach. In the near future, technologies will exist that allow integration of virtual x86 and non-x86 compute and storage capabilities. Virtual x86 and non-x86 storage and compute interoperability will become a reality with STUB technology. STUBs are similar to application API's, but focused on opening possibilities within the infrastructure OSI layer rather than the application layer. VMware's Redwood beta release includes the foundation for the virtualized infrastructure layer "STUB". The following is a description of non-x86 supportability principles that must be considered as part of the turnkey shared services design.

Legacy and non-x86 systems supportability:

COBOL is still a common language for business apps. According to a 2001 Gartner report on enterprise programming language adoption, more than 60 percent of the businesses surveyed still used COBOL. Almost every major industry relies on COBOL. Most companies that use mainframe systems for business-critical applications and data storage have continued investing in COBOL development—and probably will do so even after migrating away from mainframes. As recently as 2006, the New York Stock Exchange had moved its stock options clearing and billing systems to Windows®-based servers, but continued running its COBOL applications in recompiled form.

Over the past few years, many of the features that at one time were the exclusive domain of the mainframe have been filtering down onto lower-cost, less specialized open system (non-mainframe) hardware, based primarily on Intel and AMD processors. These commodity servers exist in a more competitive market, where innovation and development are driven hard by market forces. Scalability issues have been addressed; performance has skyrocketed, not simply at the processor level but throughout the architecture; and the harder issues of reliability, availability, and security have now been resolved.

But at the top end of the market, there are still economic and operational advantages to be gained from the mainframe platform. And in the immediate future, we see the need to continue to support those infrastructures until newer technologies in the open systems arena catch-up to the mainframe performance capabilities.

We envision in future capabilities of the proposed data center, that it will be able to make the connection between the latest generation of x86 platform services, and existing legacy and non-x86 infrastructures, through the development of programmatic "Stubs".

This integration could be made possible utilizing third generation development tools and techniques to expose legacy COBOL applications as endpoints, or merge large processing workloads to be refined and data-mined in the x 86 environments.

Mainframe and Virtual x86 Application Integration

INX

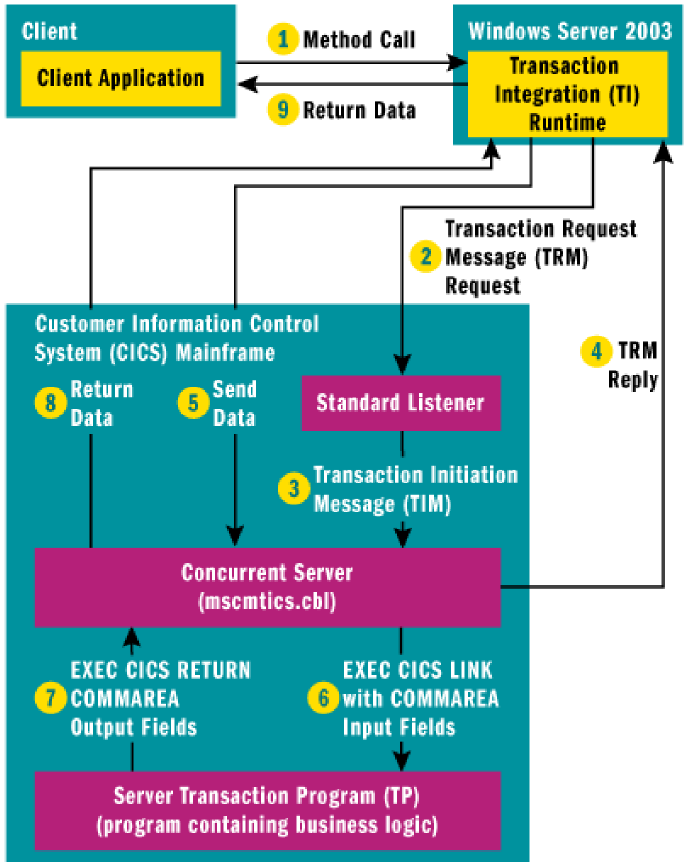

Question 8: What purchasing and technology decisions might be made now that would limit or enhance the potential for this alternative?

Network Recommendations

The proposed high level network architecture for the State Datacenter would be laid out in three functional and two services domains. The proposed end to end architecture consists of Network Core functional domain performing traffic aggregation across the LAN's, WAN's, Internet and Datacenter; Datacenter Core functional domain performing traffic aggregation between Datacenter and the rest of the network and Datacenter Access Layer functional domain providing connectivity to compute and storage platforms. In addition two services and security layers are identified; one at the Network Core domain design to provide controls for inter-network transitions (such as SGN to IGN to PGN); and one at the Datacenter Core domain designed to provide security and application services to Datacenter bound traffic.

The high level architecture will need to provide the protocol layer of the network design which is the critical layer when it comes to meeting the current needs and preparing for the future requirements. Details outlining actual traffic flows such as where layer three routing is taking place, where does MPLS tagged traffic covert to VRF and VLAN, whether security is deployed at layer two or layer three (IP layer) are critical for a successful scoping and implementation of the network of this magnitude. In addition, the high level architectures of the compute and storage platforms and virtualization design could drive up the costs and complexities of the proposed design. A detailed design of the access layer connectivity for servers must be created and is proposed as part of a holistic approach INX would employ for this project. Depending on the level of compute complexity, the type of platform chosen and the needs of the agencies on this shared service, the access layer design will be affected.

In addition, at present state, the State of WA IP Addressing Architecture (IPv4) does not allow for scaling and growth and will introduce considerable challenges during the SDC move and consolidation. It will likely be the cause of prolonged troubleshooting cycles and possible outages unless addressed. The current IP addressing architecture does not provide consolidated addressing pools required to support virtualization initiatives as well as security surrounding them. As the state moves towards shared services, this concern will become an issue and will introduce considerable roadblocks to the adoption of shared server and virtual desktop infrastructure.

Security Recommendations

Many of the security requirements needed by DIS are of a level that would have to be validated against a more detailed design which is a component of the INX proposed solution. The second is 'component specific', where a requirement is of a nature that must be tied not to an architecture but, rather, to a specific platform choice.

INX would incorporate the ESS team for collaboration in building out the security detailed design to support the states ongoing needs. The second of the two issues above is an expected constraint of the design process; the ESS team is strongly of the 'best of breed' mentality. This means that the architecture is kept as open and flexible as possible.

The lack of a more detailed design and the impact that has on validation is more severe. Before going to final implementation, a detailed design is a requirement from a security perspective and is part of the Inx proposed solution. This will cover exactly what components are required, how they will interconnect, detailed functionality and more. It is a key element that will be required in the implementation of the data center from a security perspective.

Key elements that must be clarified are how the authentication enclave will be provisioned and what its scope is (infrastructure or beyond), as well as detailed requirements for the functionality at each security checkpoint. Additionally, greater emphasis must be placed on virtualized and virtualization security, as well as on the requirements for physical vs. logical separation in the compute and storage spaces.

Storage Recommendations

DIS will benefit significantly by incorporating the backup and recovery environment for both open systems and mainframe into the solution for the new State Data Center whether managed in-house or outsourced. Backup and recovery are typically left out or forgotten in many designs. However, there are strong and compelling reasons as to why these should be incorporated, 1) DIS customers will want assurance that backups are in place to support them during and after the migration, 2) There may be instances where tape restore is the method for migrating data for a customer. A detailed design of the "end state" post migration, as well as the migration data flow and library access during the migration is necessary to ensure the solution chosen will meet the listed requirements.

Additionally, INX recommends DIS incorporate the mainframe environment more thoroughly into the design. It is expected the mainframe will be moved and supported by the warranty holder of the mainframe. The warranty holder should participate in the detailed design process to capture connectivity requirements and move constraints.

A critical component of the security design is a storage management layer. Within a detailed design, the management network, storage agents, reporting database, integration with enterprise framework, dial-in network for remote vendor support should be defined. There are several items that can be bundled into this design, one of which is customer self provisioning, as this is more than likely a function of the storage management suite of choice. So items such as customer provisioning, auditing, reporting, and mapping can be included in the storage management detailed design.

Prior to product procurement it is necessary to create a detailed design for the functions and features that are identified as requirements. Data flow from initiator to target should be mapped out, and technologies such as de-duplication should be placed within their appropriate locations within the flow. Encryption is another technology which should have its data flow within the environment identified and depicted. In short THP, de-duplication, encryption, replication, should have a separate detailed design diagram and implementation plan.

Finally, it is critical to the long term success of the datacenter and the agencies acceptance of the services available to create a detailed design that is focused on process and procedures. Capacity on demand can easily fall into this category. Another is how storage is provided to the agencies. These are more internal process flows than technology implementations but have a significant impact the solution design, configuration, implementation and future maintenance. Regardless of in-house support or outsourced completely, these process flows must be mapped out to fully understand the needs of the state and provide the appropriate level of support to satisfy those needs. DIS would benefit by creating a separate diagram specific for processes stated above.

Compute Recommendations

INX recommends DIS standardizing on a high density compute platform using a set of tightly defined protocol and access methods. INX would work collaboratively with DIS to determine the best fit to provide the services required and would then acquire assets needed to satisfy those requirements. As it is possible the outsourced datacenter will revert back to a state owned asset, INX recommends a collaborative approach to standardization and policy. A stateless platform would ease the transition of services from OB2 to the SDC by abstracting the hardware addresses. This enables portability of the operating systems so virtual machines can be moved between hosts or between datacenters without downtime. To simplify the automation within a lights out datacenter and to enable the agencies to utilize a self service portal, the management layer is critical to the overall success of the project.

INX recommends a combination of SAN and Host based replication that will be used for physical server and data migration; VMware virtual servers will utilize vMotion and Storage vMotion. Software tools will need to be identified and validated for Operating System and Application support. A more detailed design needs to be developed to identify the transport medium as well as how to move the servers and data. A detailed design is part of the INX proposed solution.

The network will provide IP and FC connectivity between the OB2 and SDC to provide the transport medium for data migration. Multiple protocols will be used during the migration including FC, FCoE, and NFS. Additionally, the storage design and implementation plan should be developed with input from the compute design as each solution set has dependencies upon each other. It is important to work with the Security Architecture team to ensure the data available meets the needs of the Security team.

Section 4 – Responses to the Financial Discussion

Briefly address the question in the Financial Section of the RFI description above. Please highlight any unique strategies or capabilities that you would provide to make this effort successful.

Question 9: What unique financial models might be necessary or advisable to maximize taxpayer value through this alternative deployment strategy?

Chargeback Model for INX recommended approach:

We envision our proposed datacenter architecture, to be capable of delivering and executing on any chargeback model required by the State.

Important recent technological advances in virtualization provide us with the unique opportunity to embed in our new architecture and design efforts, all the necessary components required to accurately track infrastructure costs in an empirical manner, rather than through the traditional approximation approach.

We now have the ability to capacity plan our projects, and dynamically estimate growth potential and trend of the individual PODs, using the new data provided to us by the infrastructure.
Concurrently, the datacenter will have the ability to mix and match between the three major traditional methods of chargeback ("Showback", "Break Even" and "For Profit") to accommodate to the individual business needs of the POD owners.

In a "Showback" model, costs associated with the different pods in the architecture can be clearly communicated to IT managers and POD owners, for the purposes of tracking and resource allocation adjustments. This would be the preferred chargeback methodology, to gradually introduce POD owners into the model without imposing on them an immediate financial burden during the initial stages of operation.

As the datacenter matures and reaches an equilibrium point between cost-efficiency and growth, a relatively more aggressive "Break Even" chargeback model can be utilized. This model would seek to recover costs associated with the implementation of the individual PODs in the shared infrastructure. POD owners would then be billed to recover costs up to a "break-even" point.
In its most mature stage, the datacenter could move into a "For Profit" chargeback model, whereas it would generate revenue from the resource consumption of the individual POD owners. Said revenue would ensure from a budgetary perspective the continued expansion of the environment, and would also provide DIS with the capital to invest in newer technology as a strategy to minimize obsolescence. This "For Profit" model could be considered if DIS chooses not to outsource the datacenter. INX understands DIS is a cost recovery model only, however for future upgrade, expansion and to keep in line with technology enhancements, DIS must consider some level of profit to support the ongoing growth of the datacenter.

This chargeback ability of the infrastructure will be capable of reporting individually on CAPEX and OPEX expenses, and differentiate between direct and indirect costs, making a huge impact in the bottom- line budgetary efforts, through the use of realistic accurate numbers as opposed to the customary approximation approach.

Although any of the three traditional cost models (“Fixed”, “Allocation-based” and “Utilization-based”) can be used in this architecture, our initial assessment is that a “Hybrid” model combining all three should be used, to accommodate for the particular business needs of each POD owner.

Initially, we would propose to track the following six (6) basic cost elements for the purposes of chargeback accountability:

Area	Capital Expenditures (CAPEX)	Considerations
Server Hardware	• Chassis, CPUs, RAM, internal disk, HBAs, NICs • Service and support costs	• Different costs depending on server performance class or configuration • Allocation vs. actual usage
Storage	• Storage hardware, SAN fabric, infrastructure, media • SAN LUNs	• Thin vs. thick virtual disk provisioning • Overhead for free space, snapshots, storage VMotion • Chargeback model may already exist • Allocation vs. actual usage
Network	• Network hardware, switches, infrastructure	• Chargeback model may already exist • Allocation vs. actual usage
Software Licenses	• vSphere, Guest OS, Server OS, database, management, automation, business continuity • Support and subscription	• Advanced vSphere features/services • Fixed vs. variable
Labor	• vSphere environment design and implementation • Provisioning, monitoring, tuning, and operating vSphere	• Time spent with management teams of dependent systems, e.g. storage and networking teams • Actual usage vs. pro-rated • Fixed vs. variable
Facilities	• Datacenter racks, power, cooling	• May already be included in pre-existing per-rack unit charges • Fixed vs. variable

Another important consideration in the design of the chargeback functionality of the proposed infrastructure is definition of the metering elements. We propose an initial set of five (5); CPU, Memory, Disk space, Disk I/O and Network I/O. It is important to note that having too many metering elements defined could severely tax the overall infrastructure resources, and in our experience, the above mentioned are the most important metrics needed to track the PODs effectively.

Fixed Cost per VM	Metering Element	Multiple Rate
Real Estate = $20	CPU – GHz Used	1.5
HA Enabled = $10	Memory – GB Used	1.8
Software = $200	Disk – GB Used	2.0
Power = $3	Disk I/O – GB used	1.4
Custom Fixed Costs	Network I/O – GB used	0

By assigning a “multiple rate” to each metering element, we will account for higher costs of storage assets vs. server assets, or distinguish between costs of development vs. production systems, and even increase rates to account for additional services such as backup and disaster recovery among others.

Non-compute elements such as software licensing, power cooling, human resources and other custom fixed costs will be tracked in custom attributes programmed in the metadata of the infrastructure.

We also envision the need to accommodate to virtualization-specific requirements that are not traditionally observed in the physical computing world:

- True cost per VM:
    - May decrease depending on host density
    - May decrease with new capabilities of the virtualization platform
        - Consider review of cost model for each generation of VMware software
- Hardware lifespan
    - VMs live longer than physical counterparts
    - OS and application refreshes can slip due to hardware independence
    - No longer tied to traditional depreciation schedule
- Overhead
    - Include host resources in addition to VM-specific needs
    - Spare capacity for FT, HA
- Service levels
    - Achieve SLAs not readily available for physical workloads

- Resource guarantees through resource pools, shares, reservations
    - Increased availability with FT, HA
    - Higher utilization and portability (VMotion)
    - Business continuity through faster replication and recovery
- Transient demand
    - Usage may fluctuate due to seasonal or cyclical demand
    - VMs can easily be furloughed when not needed
    - Give credit when VMs are not powered on

Finally, by combining all the above-mentioned chargeback methodologies and implementation techniques, our vision is to provide a fair "phone-bill-like" invoice to the POD owners. This invoice would reflect real and empirical monthly consumption as opposed to a yearly "use it or lose it" budget.

Section 5 – Cost and Schedule Estimates

The Department realizes that the proposed alternative is being presented in the most general terms. Yet there are essential elements of the model that lend themselves to reasonable cost estimates. Please provide cost information that will allow the Department to evaluate the potential for the proposed approach. Given the limited detail available for your response, respondents are welcome to provide ranges or costs rather than specific dollar estimates.

Moving from a physical to virtual infrastructure provides dramatic capital and operational cost reduction but requires a large upfront investment to begin realizing the savings. Most organizations will have to buy new high capacity servers, a storage area network, VMware licenses, support contracts, and substantial professional services and training if they are to have a successful implementation. However you can migrate from a physical to a hosted virtual server infrastructure without incurring these costs.

INX's Virtual Infrastructure Hosting Service eliminates all hardware, software, and training costs from deploying a virtual infrastructure. INX has invested in large clusters of latest generation multiprocessor servers, high performance storage arrays and switch fabrics optimized to support virtualization. By running your virtual machines in the VHS all initial and future server, storage datacenter and VMware license purchases can be eliminated.

The Virtual Infrastructure Hosting Service replaces the substantial initial investment with a fixed monthly cost almost immediately offset by the capital and operational expense savings.

The INX-proposed turnkey strategy advocates the build out of a joint asset VHS virtual infrastructure with chargeback capabilities to DIS agency customers. Below is a representative high level service catalog and cost model for network, storage, and compute cloud services. Variations of services within the cloud can be developed further as the enterprise data center service catalog is defined.

The following represents a potential cost model and timeline for joint asset virtual hosted datacenter.

Service	Cost Range	Variables Impacting Cost
Virtual Machines 5000 Virtual Machines	$75-$195/mo each VM	Number of Processors, RAM, Storage, Network bandwidth and use case
Storage 40 TB of Tier 1 Storage	$.50 - $.95/mo per GB	LUN size, performance and replication requirements
60 TB of Tier 3 Storage	$.20 - $.45/mo per GB	LUN size, performance requirements
Network 300 Mbps of Internet Bandwidth	$15 - $25 per Mbps	Geographic distance and utilization
Physical Machines 2000 Physical Servers	$95-$165/mo per rack space unit	Power requirements per servers
Virtual Desktops 1000 Virtual Desktops	$20 - $40 per vDesktop	End user application use case

INX recommends a strategy of incrementally virtualizing agencies immediately, but do so in a manner that enables virtual machines and application data to migrate over-the-wire to the SDC cloud infrastructure during the OB2 data center move timeframe. Technologies are available today that can enable this approach. The potential timeline below represents a possible sequence of events.

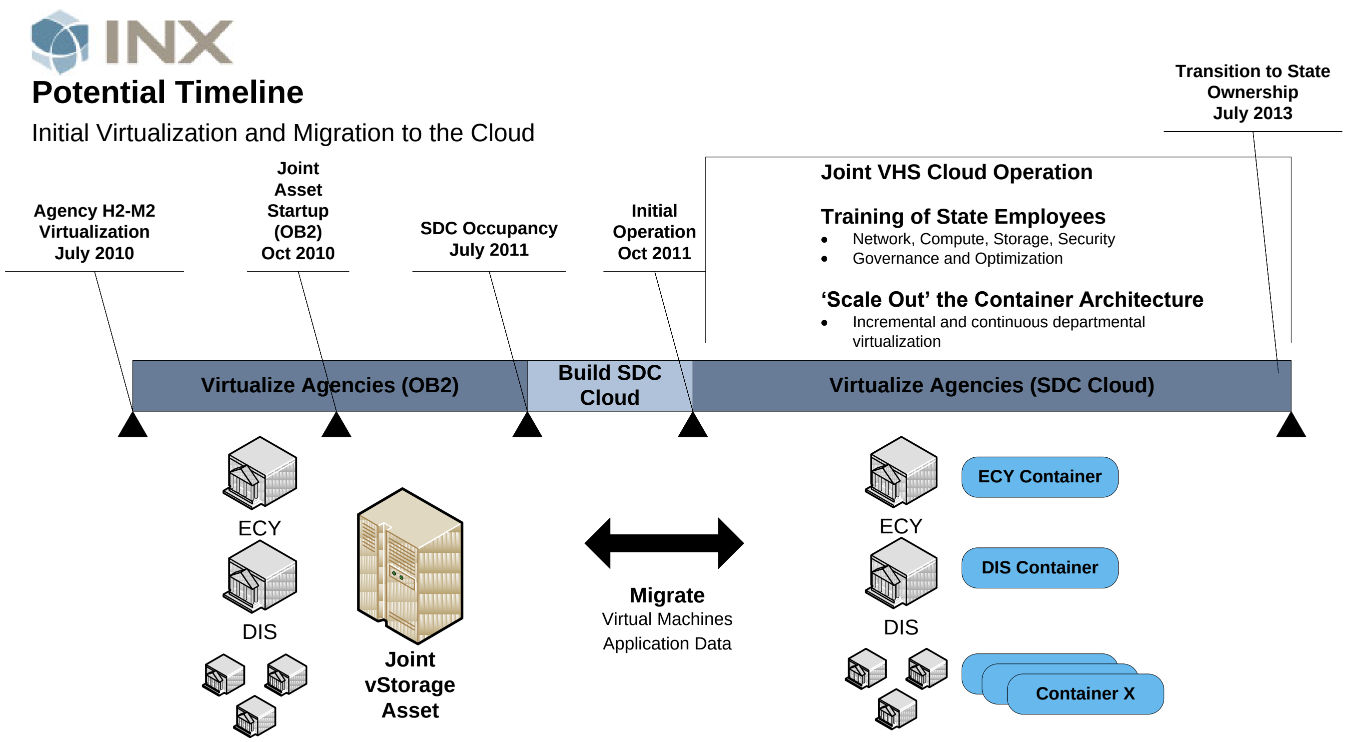

Section 6 – Corporate Expertise

Describe your company, your products and services, history, ownership, financial information, and other information you deem relevant.

In particular, please describe any projects you have been involved in that are similar in concept to what is described in this RFI, including design, management and operations approach, training provided, security assurance, relevant lessons learned, timeline to go from concept to production, and financial capacity.

Include any comments on the structure of the requirements for a formal RFP response.

INX's Strategic Delivery Framework (SDF)

The goal of the **Envisioning Phase** is to align technology usage to business needs. The Vision Team, comprised of both INX and DIS experts, is created to facilitate this process. Understanding business needs extends far beyond the function of a proposed technology into an awareness of the environment, type of industry, strategic objectives and tactical challenges of the organization. Determination of a possible solution's value to the organization is made quickly. If a potential technology application does not add positive impact, valuable time and resources are not wasted pursuing it. Envisioning is critical to the success and impact of a project. The Project Vision is the main deliverable in this phase, and is used to document the work done in the **Envisioning** Phase. Once Vision Consensus is achieved the next step is to move forward to the Proposal Phase.

The goal of the **Proposal** Phase is to refine and formalize the Project Vision. The Proposal Team, comprised of both INX and DIS experts, is created to facilitate this process. The Proposal incorporates the work done during the Envisioning Phase, and documents all key assumptions. A refinement of the proposed cost may be offered for the Execution Phase based on DIS requirements. While the Execution Phase is typically offered as a budget estimate, INX provides assurance that any remaining conjecture will be eliminated during the Planning Phase.

The goal of the **Planning** Phase is to assess and mitigate risk by documenting the project plan and creating a detailed design that considers all known issues that pertain to the technology solution, the client's business and the end users. The Planning Phase work incorporates both industry and vendor best practices as well as previous INX real world execution experience. The Planning Team, comprised of both INX and DIS experts, is created to facilitate this process. This work is centered on the discovery of the existing environment and integrating the solution into the organization. The Planning team creates the following Planning Phase documents and deliverables: High Level Architecture, System Engineering Report, Implementation Plan Scope of Work and Schedule, final Bill of Materials, and final Services Agreement to execute the solution if applicable. The work done in the Planning Phase mitigates risk by addressing all open issues and finalizing the fixed scope and budget before executing on the final decision.

The goal of the **Execution** Phase is to perform the scope of work developed and finalized in the Planning Phase. This is the phase where the actual deployment of the solution takes place. The Execution Team can be comprised of both INX and DIS.

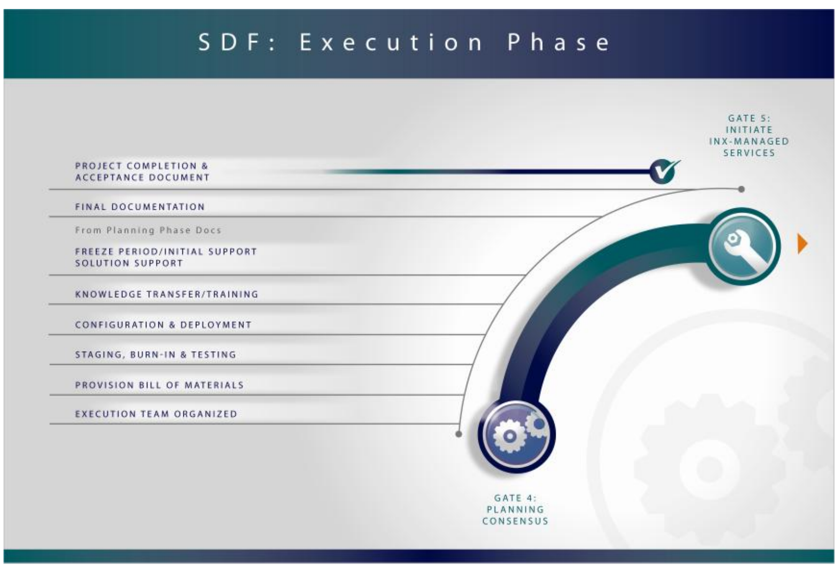

Operate Phase

The goal of the Operate Phase is to ensure a qualified, efficient, and timely support structure is in place to maintain optimal operating conditions. INX has a nationally ranked, Cisco-approved managed support service available to support the transition to the recommended technologies with the ability to provide ongoing, robust support as needed over time.

Optimize Phase

The goal of the Optimize Phase is to assure that the solution continues to add value to DIS over time. INX works with our clients on focused utilization reviews at key intervals to ensure the greatest value is being derived from the technology investments that we have worked with you to envision, plan and implement.

Approach and Methodologies

Project Management

INX provides Project Management Services to allocate and manage the resources for the successful completion of IT projects. Our Project Management Team is available to work on projects involving our technical personnel, your staff, or a combination of the two. INX's Project Management provides the expert oversight necessary to avoid delays and successfully bring together the components of even the most sophisticated and comprehensive projects by leveraging Microsoft Office Project and industry best practices methodologies (PMI) for formal Project Management.

Change Control Management

As part of managing any complex project, change is an inevitable issue. We frequently accept input from end users and assess how these change requests will affect the desired outcome. The client project lead is responsible for identifying the need for change in response to their user demand or changes in the business or project requirements. A standard change control procedure applies to all changes to provide a clear set of steps and allocate responsibilities.

Customer Service Standards and Practice

Customer satisfaction and quality delivery is our top priority. INX considers a project successful only if it is completed on time, within budget and to the customer's satisfaction. Because we use a proven and repeatable Solution Delivery Framework methodology and professional Project Management for all engagements, as well as implementing industry-standard solutions, we have achieved a very high percentage of successful Projects and repeat business from clients.

Knowledge Transfer Methodology

At INX, our definition of knowledge management and transfer refers to strategies and structures for maximizing the return on intellectual and information resources. Because intellectual capital resides both in tacit form (human education, experience and expertise) and explicit form (documents and data), INX depends on both cultural and technological processes of creation, collection, sharing, recombination and reuse.

Quality Assurance Methodology

Quality is the highest priority at INX. To that end, all INX employees follow an established, yet flexible Solutions Deployment Framework that covers each phase of project development. Quality Assurance is assured in each of the phases in the form of a formal peer review. The outcomes of these peer reviews ensure our SDF process is followed and the project is sound from a technical and functional perspective. It is this strict adherence to quality assurance and the SDF process that ensures better business results with lower risk.

Our methodology begins with the requirements gathering phase. This process is carefully managed by the INX Systems Engineer and includes an INX Project Manager working closely with your company. Once the project requirements are finalized our quality assurance team maps the requirements document to the specifications to ensure that all functionality is addressed and that the project we've specified is exactly what you want.

Upon completion of project implementation, INX uses a formalized Test & Acceptance plan prior to final delivery and sign off by the client. The intent is to validate with the customer that the system functions as contracted.

Accelerating the Adoption of virtualization

Many enterprise organizations are accepting the benefits of consolidation and centralization of services. There are several examples where consolidation projects were undertaken only to learn far too late in the process that critical steps had been missed and a narrow view of the future of the consolidated datacenter limited the financial savings and actually substantially increased the cost of the new datacenter.

INX's team has successfully completed numerous data center consolidation projects and developed the specialized skills and processes that ensure success. From developing and executing the RFP bid process and site selection, to envisioning, planning, designing and implementing datacenter infrastructure including virtualization of servers, storage and networks, efficient power, cooling, network cabling, rack layout and modular-based cabinet designs. Our team has worked closely with our clients through all phases of end-to-end project lifecycle.

Our Solutions Delivery Framework methodology ensures a consistent approach to projects regardless of technology. Each phase of the SDF considers the client stakeholders most accountable for that specific phase. Within the ROI/TCO phase of a project the lead client stakeholder might be the CFO and CIO, whereas the core network design lead stake holder might include both of those individuals as well as the client network engineers. This approach ensures the specific requirement of each group within the client organization has input into the project and a vested interest in the successful outcome of the project.

Within the SDF methodology, the **Design, Build, Migrate** phase is conducted in a team approach. Each successful prior engagement has had a senior data center project manager leading the engagement, and supported by data center architect(s) as well as subject matter experts in the infrastructure, network, server and storage disciplines.

INX Virtual Hosting Service®

INX has in production 2 tier 4 datacenters that are 60 Kilometers apart that provide the Virtual Hosting and Virtual Recovery services we provide to our customers. These two datacenters, while geographically separate appear as one logical datacenter to our practice and perform as a clustered datacenter. This allows our Virtual Hosting Service (VHS) and Virtual Recovery Service (VRS) customers to seamlessly move VM's between datacenters with no interruption of service and no detectable change to their end users. Currently INX supports more than 50 customers and 7000 virtual machines. Additionally, INX maintains multiple additional datacenters across the US that provides Disaster Recover and Business Continuity to overall national datacenter practice. INX's multi-million dollar investment in VMware ESX servers, enterprise SAN storage, and a redundant, high-speed network infrastructure, offers organizations one of the finest cloud computing experiences available in the market today. The cornerstone of INX's turnkey approach leverages our virtual hosing business lessons learned and the accompanying mature cloud architecture.

INX Virtual Recovery Service®

INX's Virtual Infrastructure Recovery Service (VRS) includes a comprehensive virtual infrastructure based disaster recovery plan. The Virtual Infrastructure Recovery Plan looks at all critical information processing area of the company, including but not limited to the following:

- LANs, WANs, and servers
- Telecommunications and data communications links
- Applications, software, and data
- Data backup solutions including disk and compression
- Recovery methods
- Workstations and workspaces
- Scalable and highly available cloud service offerings

INX has extensive experience in broader based disaster recovery / business continuity including working with small, medium, and large organizations. We have written about different aspects of disaster recovery in multiple books and articles, and have helped virtualize companies of all sizes. VHS/VRS offers a world class designed and engineered N+1 (minimum) co-location facility with security, uninterruptible power, totally redundant, carrier neutral fiber networks, state of the art environmental and fire detection/suppression systems, and is SAS 70 Certified. VHS/VRS is an architected as a 99.999% mission-critical, co-lo data center with 24 x 7 staffing and client access.

INX Hosted Virtual Desktop Case Study (currently in Implementation Phase – near completion)

INX recently transitioned a CPA firm from traditional client server desktop network architecture to a 100% virtually hosted solution including all physical desktops. The CPA firm has three offices in three separate cities and provides a diverse range of auditing, accounting, tax, and recruiting services. Any interruption to productivity is very costly to the firm. Also contributing to a significant administrative overhead is the constant upgrading of desktop software (accounting and tax) as new versions are released. In a traditional environment each new release would require that the software be installed on each workstation separately. Because of the amount of software applications installed and the remote connectivity requirements the firm found itself both having to spend more internal time on the network as well as the additional cost of having to contract with a technology partner for most major upgrades and troubleshooting.

In summary the major requirements and goals of this project included:

- Leverage pre-existing virtual infrastructure
- Remove the need for reoccurring hardware upgrades (both desktop and server)
- Allow for access to desktop regardless of employee location
- Significantly reduce in-house IT administration
- Eliminate the need to contract out projects to third parties
- Allow for rapid desktop provisioning
- Have a desktop support structure that included 24X7 support and access

INX's solution to address these needs included:

- 100% hosted servers (AD, Database, Exchange, and CRM) using VMware
- 100% hosted desktops using VMware View 4
- VRS services for disaster recovery and data backup
- 24X7 monitoring for network infrastructure and desktops
- 24X7 support desk for infrastructure and application support

ORGANIZATIONAL CAPABILITIES

At INX, our product is not technology, but professional consulting services wrapped around technology to create solutions that meet our client's specific needs. We envision with our clients in a collaborative team setting to devise the solutions that can most effectively enable their organizations and the business processes that run them. The collaboration team is comprised of consulting specialists experienced in business applications, work flow, and operational interdependencies along with technology subject matter experts and the clients own stakeholders that have expertise in the organizations form, function and expected outcomes. This collaboration team envisions and begins to architect the high level solutions that meet the requirements identified. This iterative and collaborative team approach ensures that the solutions are purpose built to specifically meet business goals. By aligning the business objectives with the technology capabilities we can ensure the solutions hit the mark whether they are improved communications and increased collaboration, shared services, reduced operating costs, and/or better productivity.

In 2009 INX was awarded the **Americas Solution Provider of the Year Award at VMware Partner Exchange**. Furthermore, in 2009 INX was awarded the **Cisco Systems Global Technology Excellence Partner of the Year for Data Center** and in 2010 the **Cisco Systems Services Partner of the Year**. These distinctions are the highest honors bestowed by Cisco out of 60,000 worldwide partners. INX has the same level of commitment to our storage and virtualization partnerships and as such have achieved the highest level of certifications as a Star Partner with NetApp, a Star Partner with EMC, and a Microsoft Gold Partner. These certifications place INX in the top tier of solution providers and are the result of exhaustive audit processes by each of the manufacturers to evaluate our business process, delivery methodology, engineering certification, sales certification and expertise as well as customer satisfaction.

Our corporate headquarters is in Lewisville, TX with 10 regions nationally, internationally and growing. Our employee count is 450+ with more than 130 advanced technology consultants. There are 5 offices in the Pacific Northwest; Portland, Eugene, Boise, Seattle (Bellevue) and Olympia with more than 60 employees. For over 25 years, INX has been refining internal processes to create a deployment methodology (SDF) that ensures IT projects are completed on time and within budget.

We aspire to be a part of something bigger than technology alone: we are agents of change, innovation and thought leadership. We transform the use of information by collaborating with clients and empowering them to leverage technology for more efficient business operations. Our identity is rooted in our engagements with our clients and tied to their success at every level. Our clients are not just an extension of our business; they are our business as we engage with them from envisioning to execution. Our approach to advanced technology, holistic and comprehensive in its blend of people, process and technology creates lower risk for the client in a way that provides significant value differentiation. Our years of collaboration between our practice/solution development teams and our delivery methodology have resulted in the industry's most reliable, proven solutions.

Virtualization Project Examples

The following is but a small sampling of some of INX's many DC consolidations and implementation planning projects to demonstrate our understanding of the unique customer requirements in designing, building and migrating a datacenter.

ProLogis Project title: Datacenter Move

INX was contracted to move the ProLogis data center from Albuquerque, NM to Denver, CO. This project consisted of a single main datacenter, several 'secondary' datacenters and many smaller 'server rooms' spread throughout the United States. In addition to the datacenter move, INX developed a Business Continuity Plan (BCP) for ProLogis to mitigate the risk of loss of business during the actual move. In the event of failure during relocation, this plan provided specific direction as a guide for recovery. This plan was used to advise ProLogis of tasks and actions to be completed prior to the relocation for a preemptive and proactive approach for relocation success.

Overview of Project Deliverables:
ProLogis had 32 sites to consolidate server and data into one main data center. Each remote site had direct attached storage with a variety of software protocols, line of business applications and back up strategies. All sites needed to replicate to a remote location for Disaster Recovery and Business Continuance policies. Two-way replication was a requirement of this project as remote offices and the main DR site all had to have a mirror image of the data. The PeopleSoft /Oracle environment required high availability and included Microsoft SQL and Exchange while uptime was required to be at 99.999%. For the existing main site, the destination site and the 32 remote sites, the following tasks were performed to ensure overall project success.

Redundancy – Mission critical applications data was replicated to an alternate facility. A pre-positioning plan was implemented to mitigate the risks. ProLogis elected to perform the data replication in-house and was allotted three weeks to ensure redundancy and smooth cutover.

Data Protection – INX testing included "Snap Mirroring" of the data to the Denver location and verification of that data was checked regularly. Complete backups of the data were done to provide an additional level of protection.

Back out Plan: If equipment does not arrive –it was necessary to ensure that adequate redundant and/or replacement temporary equipment was available to resume business operations. INX secured virtual devices to ensure mission critical applications could be restarted.

Getty Images Project title: Datacenter Plan and Move

INX was contracted to plan the consolidation of Getty Images' seven data centers and move their main data center from Seattle, WA. This project consisted of a building out a new 6,000 sf datacenter and moving their equipment and applications. Their computational environment included 1,000+ virtual and physical servers, 1.5 petabytes of EMC storage, mainframe, tape backup and Cisco network.

System Migration – Mission critical application data was replicated to the new facility utilizing seed equipment and kept in synchronization until the system was cut over. The application servers, both virtual and physical, were staged in the new data center prior to the system migration. For each complete system (applications, data, middleware) the applications were migrated electronically over the network backbone, brought up and tested before swinging them into production.

Transporting the Equipment – Equipment such as EMC storage and IBM mainframe which were under maintenance agreements were scheduled and managed using the maintenance vendor to ensure continuation of coverage. 12 separate move events were required to maintain business continuity and minimize operational impact.

Back out Plan: Each system migration and equipment move had a roll back plan with multiple go/no go steps. Each step had minute by minute times to ensure activities occurred within the allotted maintenance windows. All of the systems and equipment were moved without the business or customers negatively impacted.

Infospace Project title: Datacenter Go Live

This project consisted of a building out of two new data centers, deploying network, server and storage equipment followed by migrating services and cutting them on-line. Infospace's business required that the equipment, applications and data move without disrupting their 7x24 service.

Infospace uses a number of technologies to deliver world class metasearch services to their customers. This technology includes: NetApp storage, Cisco, Checkpoint, Netscaler, HP Storage, HP Blades, Netezza, VMware, Linux, MS Windows 200X, MS SQL, Oracle, MS Exchange, Active Directory, SAP.

System Migration – Mission critical application data was replicated to the new facility and kept in synchronization until the system was cut over. Because of the complexities and interrelationships of the applications and data, this move required a complete data center cutover. In order to remove risk, as each new data center was brought online, traffic loads were managed and monitored to ensure quality.

Transporting the Equipment –Transportation of equipment was managed through procurement and the transportation vendor for the legacy equipment to ensure proper precautions were taken when moving sensitive electronic equipment. Separate move events were required for each data center to maintain business continuity and minimize operational impact.

Back out Plan: Each migration had a roll back plan with multiple go/no go steps. Each step had minute by minute times to ensure activities occurred within the allotted time frames. All of the systems and data were moved without the customers experiencing any down time.

Banner Health: DC Virtualization with 850 servers and virtualization of 2,800 PC's

INX's work directly resulted in an enterprise-wide virtualization initiative including a 2,800 desktop virtualization pilot. Also included were VMware automation products: Site Recovery Manager and LifeCycle Manager. INX worked closely with the VMware and Banner IT staff teams on project planning.

INX met with the Banner CIO on April 22, 2008 to present the findings of the Capacity Planner assessment and to present an initial high-level ROI analysis. The CIO was extremely enthusiastic about the results and made the statement, "If we give you the go-ahead now, how long will it take to virtualize us?"

Seattle University: Centralized computing and virtual desktop
(600 Virtual Desktops in 7 computer labs and 100 classrooms)

The centralized computing and virtual desktop project leverage leading edge technology of the next-generation data center to transform the end user experience for students across the SU campus. The solution consolidates resources in the data center, creating efficiencies and increased availability of applications while driving down costs such as power and cooling. Built upon the Cisco Systems Unified

Computing System (UCS), INX deployed a solution that would transform SU's desktop computing model by enabling secure central desktop computing for their users.

In addition to centralization, a virtual desktop infrastructure (VDI) was implemented which moves all of the applications of faculty and students off of their individual computers and into the data center. This technology allows the computers to be simplified terminals or clients that receive the applications and content which is hosted and managed more efficiently in the data center. The VMware VDI solution will simplify SU's desktop management strategy while increasing control of University assets and significantly improving the security of end-user data. Students will have ubiquitous wireless and wired access on campus. Additional third party applications have been incorporated to allow for single sign-on to the desktop as well as printer location management. By virtualizing the desktops, SU will be able to address the requirements for throughput, redundancy, manageability, and flexibility.

Clean Water Services Project Title: Virtual Desktop Infrastructure Proof of Concept

A Proof of Concept (POC) with a small test set of virtual PC's was evaluated. INX designed the solution to function independently of the hardware and operating system to provide the client devices with a broad platform choice (thick/thin/zero client). As a result, the solutions provide a key integration point for hardware and infrastructure management vendors to deliver differentiated value that can be applied uniformly across all application and operating system environments.

With VMware VDI, end users received a complete, unmodified virtual desktop that behaved just like a physical PC. There was no change to the applications or desktop environment, no application sharing and no retraining required. Administrators could decide to allow users to install applications, customize their desktop environment and use local printers and USB devices. However, lockdown of the virtual desktop was also an option.

Middlesex County Project title: Datacenter Consolidation to include virtual desktop

Middlesex County had an initiative to implement a consolidated datacenter and secondary disaster recovery solution for its 50+ production servers at 14 different sites. The solution was to virtualize and consolidate the production data centers as well as set up an automated virtualized DR facility.

INX also designed the utilization of VMware's new data center automation products in order to manage processes such as workflow approval for provisioning new servers, and staging the virtual machines before they are put into production.

Subaru of NE Project title: New Datacenter to include DR & BC

INX performed a virtualization Capacity Planner assessment and corresponding ROI analysis to demonstrate to Subaru NE that by its production data center not only would it gain benefits in both high availability and performance, but would also reduce its IT budget in the process.

Subaru of NE is currently in the Planning Phase for a Virtual Desktop Infrastructure implementation project utilizing VMware View 3 and Pano Logic zero client devices. This project is building upon the success of the datacenter virtualization project and incorporates the significant ROI and TCO analysis completed by INX to support the business justification to embark on this new technology. Subaru of NE believes the long term financial benefits of reduction in operating expenses to support a virtual desktop will help them maintain the edge in their competitive environment.

Walnut Creek, CA Project title: Multi-phased Datacenter Virtualization with energy rebates.

Walnut Creek awarded INX a contract to design and implement a multiphase virtualization solution that included a server consolidation effort, a backup and recovery solution, and a virtual desktop initiative. The project leverages virtualization technology from VMware to create efficiencies in data management thereby decreasing the amount of servers required and increasing the reliability of application and data delivery. The virtualization technology also enables the city to eliminate traditional desktop computers and replace them with low energy consuming terminals. The applications move to the data center where they are centrally managed and deployed, which drives down costs, while enhancing the user computing experience.

The City of Walnut Creek is currently in the evaluation stage of implementing a virtual desktop infrastructure. ROI and TCO analysis support a long term business justification and the next step will be a pilot project for proof of operability with their current applications.

Experience of Staff

INX maintains a staff of highly experienced and certified consultants and engineering resources to envision and deploy our client's technology project. All INX staff members recommended below for this project meet the Mandatory Experience Qualifications as outlined in the Work Order. As an attachment to this RFP response we have included the resumes of the following INX resources available for this project:

Cary Dahl, Sr. Virtualization Consultant: Solutions Architect with extensive experience in building scalable IT infrastructure for high-tech businesses. Product knowledge spans telecommunications, facilities, IT infrastructure, and Decision Support Data Management. Related areas of expertise are:

- Storage Area Network Design and Deployment
- WAN/LAN Architecture
- Virtualization – Application and Server Using VMware Infrastructure
- Storage/Server Consolidation
- Microsoft Active Directory/Exchange
- Structured cabling and network segmentation
- Former Seattle VMware User Group President

Julian Zbogar-Smith, Sr. Messaging & Virtualization Specialist: Experienced Exchange professional with 7+ years working with all versions Exchange (5.5, 2000, 2003, 2007, 2010 beta) Julian has broad knowledge of enterprise Microsoft products, but happily admits that Exchange and Active Directory are his favorite solutions and enjoys that specialization due to its complexity, dependency on many key areas of enterprise infrastructure, and critical business value. He is proficient in the design, implementation, troubleshooting, and support of Microsoft Active Directory enterprise networks, and a variety of other protocols and technologies

Al Kari, Director of Consulting Services: Driven, innovative and result-oriented executive with a track record of success in critical strategic planning and tactical execution. Proven results in forming, leading, growing and developing high-performance teams, achieving business objectives and driving growth while delivering world class service to Fortune 500 companies. Currently focused on blending business needs with data center optimization, server consolidation and virtualization technologies. Al is a creative communicator, and able to work with customers at the CXO level.

Gary Lamb, Sr. Director of Data Center Virtualization: Gary Lamb is a nationally known expert in the virtualization field and is one of 30 members from around the world on the VMware Partner Advisory Council. Gary's presentation at Virtualization Congress 2009 was voted one of the most five popular sessions. Gary was a co-author of the Sage book, Deploying the VMware Infrastructure, and has co-authored several virtualization articles and white papers. With 27 years in IT, Gary has worked as a MVS Systems programmer, LAN/WAN design and integration architect and an enterprise security assessment practitioner, including running a consulting practice for 2 ½ years for ISS. He also has an extensive background in disaster recovery. Gary has helped several manufacturers build their products to better interoperate with a virtual infrastructure.

Ed Gruss, Director of Technology Infrastructure: Senior Systems Engineer with 20+ years of comprehensive experience in the network industry. He is proficient in the design, implementation, troubleshooting, and support of enterprise networks including LANs, WANs, IP Telephony, network security, firewalls, VPN/IPSEC, VLANs, Frame Relay, ATM, DSL, TCP/IP, IPX/SPX,SONET and a variety of other protocols and technologies from a carrier class perspective. As the Director of Technology and Infrastructure, Ed is instrumental in working with INX clients to design and configure a strategic virtual infrastructure network which optimizes performance and return on investment.

Joe Gagznos is an INX Sr. consulting engineer and experienced technology professional with over 8 years of experience in the networking industry. He is proficient in the design, implementation, troubleshooting, and support of enterprise networks including LANs, WANs, firewalls, VPNs & IPSEC technologies, VLANs, Frame Relay, xDSL, QoS, OSPF, BGP, and several other technologies. Joe is also experienced with VMware and Microsoft enterprise software.
Recently, Joe was the lead networking engineer involved with the $1.75M State of WA – Dept. of Ecology Core Infrastructure Upgrade project involving the Cisco Nexus 7K, 5K, 2K and Nexus 1000v platforms. This was a highly successful project completed in December 2009 and was completed on time and within budget. Consequently, INX is the only Cisco fulfillment partner on the State of WA Cisco Master Contract that has the required ATP certifications to sell and implement Cisco Nexus equipment.

Goran Ognjenovic is a dual Cisco CCIE (Service Provider and Routing/Switching) and Sr. INX Consulting Engineer with significant experience in the design and implementation of large MPLS, LAN/WAN, Cisco IP Telephony projects, Cisco security projects, LAN and migration strategies. Utilized technologies based on variety of Cisco hardware and software platforms including MPLS, Cisco Call Manager, Cisco Unity, Cisco 2600, 3600, 7200, Cisco Catalyst 4006, 4507, 6509 and Cisco Aironet wireless.

Goran has extensive experience in architectural designs and implementation of large (service provider) networks. He has developed and implemented high complexity BGP networks, as well as a global, multi-carrier MPLS network based solely on Cisco architecture and best practices.

Charlie Winckless is an INX Security Solutions Principal Consultant with 15+ years of progressive experience in the networking industry. Charlie currently holds the prestigious CCIE certification with a Security Specialization. He is proficient in the design, implementation, troubleshooting, and support of enterprise networks including LANs, WANs, IP Telephony, network security, firewalls, VPN/IPSEC, VLANs, Frame Relay, ATM, xDSL, QoS, OSPF, BGP, and a variety of other protocols and technologies. Charlie also has experience with web design, email, and DNS configurations.

Charlie Winckless has worked with multiple clients of all sizes, advising them on the design and architecture of their security. He has also achieved non-vendor specific industry certifications such as the CISSP and CISM further validating his knowledge of these areas of security.

Michael Calderon, Virtualization and Cloud Computing Specialist: With more than 30+ years of experience in architecting IT solutions, Michael has contributed to the assessment, planning, design and deployment of IT technologies to meet business objectives. Mike's consultative approach and technical knowledge spans multiple areas of the IT infrastructure including Applications, Server Platforms, Operating Systems, Networking Infrastructure, Desktops and Storage components. Mike's skill set are applied to production deployment of applications encompassing business continuance/disaster recovery issues & operational concerns with a focus on cost-effective solutions to meet the strategic goals.

Steve Kaplan, National Director of our Datacenter Virtualization Practice: Steve is currently one of only 300 VMware Experts in the world. Formerly a Microsoft MVP, Steve co-authored 8 Official Citrix Guides in conjunction with Citrix and co-authored an 'Introduction to Virtualization' and 'Foundations for Cloud Computing' (book covers below) in collaboration with VMware.

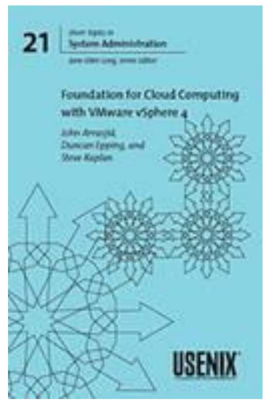

Additional Materials and Others Items We Should Consider

Please provide any other materials, suggestions, and discussion you deem appropriate.

Virtual Infrastructure Hosting Service (VHS)

According to Gartner and Meta Groups, one of the fastest growing trends in the IT Industry is the movement toward hosted, managed systems. Organizations are realizing that in the world of Windows and Linux, there is no longer a strategic advantage in owning and operating their hardware and maintaining company owned data center/server rooms. With the cost of hardware having reached commodity pricing, reducing the high cost of systems and network administration is a key to reducing the cost of IT. Data Center/Server Rooms are more densely populated leading to increasingly costly expansion projects and higher power and cooling costs.

With the advent of Data Center Virtualization, INX Managed Services has developed our Virtual Infrastructure Hosting Service. INX's multi-million dollar investment in VMware ESX servers, enterprise SAN storage, and a redundant, high-speed network infrastructure, offers organizations one of the finest cloud computing experiences available in the market today. VHS offers organizations the following benefits:

- **Located in multiple, highly secure Tier IV data centers with 99.999% very-high availability**
- **No upfront or costly capital expenses for servers, SAN storage and network infrastructure**
- **Reduced operating costs for systems administration and systems maintenance**
- **INX architects and engineer employ industry best practices to design and operate the VHS**
- **INX periodically upgrades the VHS – so you are operating your business on the newest and fastest computing infrastructure**
- **Physical Systems often involve the high cost of colocation**
- **VHS is monitored and managed 7x24 by fully trained professional engineers and technicians**
- **VHS has been hosting production and e-commerce web sites for over 3 years**

Hosted, Managed IT Systems (Cloud Computing) offers organizations the flexibility and performance they need to operate in today's rapidly changing environments - helping them harness the power of the internet

How Does VHS Work?

Building on INX's knowledge and expertise in virtualization, networking, storage and the data center, the Virtual Infrastructure Hosting Service ® (VHS) delivers unsurpassed highly-available capabilities utilizing VMware ESX and other VMware products, monitoring and management software, multiple, fiber channel SANs, a combination of the industry's newest 4-processor, quad core servers, Gig-E switches, firewall, load balancing and other components. It includes 24 X 7 monitoring and management along with BCP connections to multiple fiber carriers. The service is hosted at one of INX's multiple, top tier data centers,

VIRTUAL HOSTING SERVICE (VHS)

with minimum N+1 redundancy in all electrical, cooling and fiber network systems – data centers designed, built and operated for 99.999% availability.

INX's onsite engineers and technicians are fully trained and experienced in all aspects of the VHS hardware and software. Using VMware's Virtual Infrastructure as the foundation, VHS is architected, built and managed to achieve 99.95% system availability including both - planned and unplanned downtime. INX engineers work closely with clients to quickly handle problems that occur, i.e., such as restoring accidently deleted files, adding additional memory or storage capacity when needed.

Virtualize Disaster Recovery!

Using INX's Virtual Infrastructure Hosting Service ®™ (VHS), allows for the simple implementation of INX's Virtual Infrastructure Recovery Service ®™ (VRS) which enables failover of all X86 servers to our second data center in the event of a disaster. A virtual replica of your production environment is up and running in less than 4 hours! Because we virtualize client access, we enable continued computer operations in the event of a prolonged pandemic disaster.

Why INX?

A VMware Premier Partner, INX has extensive experience in broader based disaster recovery / business continuity including work with large financial organizations. Our pioneering efforts in virtualizing disaster recovery have been recognized in the form of speaking engagements around the globe including *Continuity Insights Management Conference*, *VMworld* and *ThinPower*. We have written about different aspects of disaster recovery in multiple books and articles, have a patent pending for our methodologies and have helped virtualize organizations of all sizes.

INX offers a world class designed and engineered N+1 (minimum) co-location facilities. With the finest in physical and electronic security, uninterruptible power, totally redundant, carrier neutral fiber networks, state of the art environmental and fire detection/suppression systems, and SAS 70 Certified; INX data centers are Tier 4, Class A rated (both the highest accredited levels) 99.999% mission-critical, co-lo data center with 24 x 7 staffing and client access.

Cost

Cost for the Virtual Infrastructure Hosting Service® depends upon the specific configuration and includes significant volume discounts. The price includes the cost savings normally associated with the server consolidation benefits of virtualization. The ROI for using VHS provides a compelling case for the use of this great service.

INX MANAGED SERVICES

vmware® premier PARTNER

Virtual Infrastructure Recovery Service (VRS)

Research conducted by Meta Group shows that most organizations suffering a disaster, such as a fire or flood, are out of business within two years. Gartner says that 40% of companies that experience a major disaster will go out of business if they cannot get access to their data within 24 hours.

The tight coupling of Wintel-based server hardware with the operating system makes recovery very costly and challenging.

- **Physical Systems most often fail to Recovery per Plan**
- **Physical Systems are costly and require twice the H/W & S/W Investment**
- **Physical Systems require double the amount of Administration time and effort**
- **Physical Systems often involve the high cost of colocation**

"Statistics show that 96% of companies that lose access to their data centers for 10 days or longer are out of business within a year!"

Virtualize Disaster Recovery!

INX's Virtual Infrastructure Recovery Service ®™ (VRS) enables failover of all X86 servers to our data center in the event of a disaster. A virtual replica of your production environment is up and running in less than 4 hours! And, while not required, virtualizing your production environment will produce savings that more than offset the cost of the DR solution. Because we virtualize client access, we enable continued computer operations in the event of a prolonged pandemic disaster.

How Does VRS Work?

Building on INX's knowledge and expertise in virtualization, networking, storage, replication and the data center, the Virtual Infrastructure Recovery Service ® (VRS) delivers unsurpassed recovery capabilities utilizing VMware ESX and other VMware products, replication methodologies, monitoring software, multiple fiber channel SANs, a combination of the industry's newest 4-processor, quad core servers, Gig-E switches, firewall, load balancing and other components. It includes 24 X 7 monitoring and the ability to instantly turn up the Internet connectivity to whatever level is required for remote access in the event of a disaster. The service is hosted at one of INX's multiple, top tier data centers, with minimum N+1 redundancy in all electrical, cooling and fiber network systems – designed, built and operated for 99.999% availability.

Virtual Infrastructure Recovery Plan

INX's Virtual Infrastructure Recovery Service ® (VRS) includes a comprehensive virtual infrastructure based disaster recovery plan suitable for presentation to the Board and to auditors. It provides an organized way to make decisions if a disruptive event occurs. The Virtual Infrastructure Recovery Plan looks at all critical information processing areas of the company, including but not limited to the following:

VIRTUAL RECOVERY SERVICE (VRS)

- RPO/RTO
- Telecommunications and data communications links
- User access / authentication
- Applications, software, and data
- Storage, backup and replication
- LANs, WANs, and servers

The Virtual Infrastructure Recovery Plan allows us to tailor the Virtual Infrastructure Recovery Service® (VRS) for your environment including factors such as quarterly testing parameters, user support, failover process and future scalability.

Why INX?

A VMware Premier Partner, INX has extensive experience in broader based disaster recovery / business continuity including work with large financial organizations. Our pioneering efforts in virtualizing disaster recovery have been recognized in the form of speaking engagements around the globe including *Continuity Insights Management Conference*, *VMworld* and *ThinPower*. We have written about different aspects of disaster recovery in multiple books and articles, have a patent pending for our methodologies and have helped virtualize organizations of all sizes.

INX offers a world class designed and engineered N+1 (minimum) co-location facilities. With the finest in physical and electronic security, uninterruptible power, totally redundant, carrier neutral fiber networks, state of the art environmental and fire detection/suppression systems, and SAS 70 Certified; INX data centers are Tier 4, Class A rated (both the highest accredited levels) 99.999% mission-critical, co-lo data center with 24 x 7 staffing and client access.

Cost

Cost for the Virtual Infrastructure Recovery Service® depends upon the specific environment, but is exceptionally affordable even without considering the huge savings available from virtualizing the production environment. Cost for the Virtual Infrastructure Recovery Plan depends upon the size and complexity of the operation and the services desired. It is almost always less costly than building a DR system and collocating it.

INX MANAGED SERVICES

vmware® | premier PARTNER